Kodak Graphic Communications Canada Company
4225 Kincaid Street
Burnaby, B.C., Canada
V5G 4P5
Tel: 1.604.451.2700
Fax: 1.604.437.9891

Request support:
ecentral.creo.com
http://graphics.kodak.com

Internal: 710-00211A-EN-A
Release date: 20 Aug 2010

White Paper

Tim Adams
Technology Infrastructure Manager
Tel: +1.604.451.2700
tim.adams@kodak.com

Virtualized Systems as a Basis for Redundancy

Foundation for Business Continuity, Disaster Recovery, and Managed Failover

Many people seem to know they want or need to virtualize their systems, but far less seem to have a good grasp of why they should do so. Virtualization holds a great deal of promise for redundancy to support business continuity and similar needs; the big question is how much capability comes “out of the box” and how much other work is required to have a robust solution?

This document will discuss several levels of infrastructure redundancy with a view to accomplishing maximum possible uptime based on various risk factors, and suggest several possible architectures that correspond to different levels of risk mitigation. These will range from internal factors within a given component’s design to large-scale infrastructure designs to mitigate catastrophic failure of a site. Several of these factors are not unique to a virtualized environment and will be identified. While several different parts of the overall infrastructure will be addressed, areas such as power and Internet Service Provider (ISP) redundancy require a far more extensive investigation and this paper will only touch on the topics needing to be covered.

Types of Redundancy within Individual Components (Server and Storage)

Within a given host computer, it is recommended that you have at least two forms of redundancy, assuming there is no local storage. Local storage, that is disks or RAID arrays installed within a single server, are not available to other servers in case of system failure and are therefore not suitable. In addition, storage that is considered highly available requires some specific design features for redundancy.

- **Power Supply Redundancy** allows you to continue full operation in the event of a single power supply component failure, along with the ability to hot-swap a replacement power supply for the failed unit without interrupting production. Within Kodak, any of the Premium Workflow servers come with hot-swappable redundant power supplies by default. If selecting host computers from another vendor, ensure that the systems contains redundant power supplies. This also applies to any server or storage component in the system.
- **Network Interface Redundancy** provides the ability to continue full operation in the event of a single port or network interface adapter failure. This applies to servers and storage components. The rack-mount premium workflow server from Kodak provides two pairs of teamable network interfaces, which allows you to aggregate throughput using a virtual IP address and transparently failback to single

interface throughput in the case of a port failure. Any storage components should also have redundant interfaces, whether network (such as iSCSI) or other (such as fiber channel).

- **Storage and Storage Controller Redundancy** is a requirement for highly available storage, and any SAN or NAS storage must provide not only redundant disks (such as RAID configuration) but also multiple paths of access to data through redundant storage controllers, or RAID controllers. Systems like Isilon have sufficiently virtualized storage arrays that they provide equivalent support to redundant controllers, and systems like Dell's Equallogic SANs provide dual controllers for redundancy, with each controller supporting its own redundant network interfaces.

A further note about storage is that virtualized systems operate very efficiently with block-based storage such as an iSCSI Storage Area Network (SAN), where the storage file system updates only small parts of larger files that have been changed. Prinergy usually deals with unstructured file-based storage and most of the time requires complete files to be written; as a result a Network Attached Storage system that provides share-based file serving often provides best performance for a Prinergy multi-user environment. For virtualization and job data sharing, a solution that provides both block-level access and file-level access is ideal. Network performance is critical in a virtualized environment running on network-based storage, and usually VM-related traffic is segregated from client network traffic for performance.

Types of Redundancy within the Infrastructure

- **Host Redundancy** is providing sufficient computing power (network, processor, and memory) to support clients, or move or restart systems in case of the failure of another host. Host redundancy is supportable in several ways in the virtualized workflow environment:
    - Emergency Plate Making (EPM) for Prinergy Connect is supported within the virtualized environment. This allows you to have a second server with a synchronized database that can be promoted to active primary in the case of the failure of the original primary. EPM can be supported within virtualized systems, and between physical and virtual systems using the same operating system. This is valuable redundancy at the operating system level, as it provides a second copy of the database in the case of corruption of a single virtual machine.
    - InSite Prepress Portal supports a configuration with two active servers where both servers are licensed. This configuration requires minimal configuration changes to support the failover. InSite Storefront is not supported in the model.
    - High Availability in VMWare's VSphere allows for a virtual server to be restarted on a new host computer in the event of host failure. This failover strategy may require custom development work depending on the requirements of the customer. The user experience does include potentially seeing an error message but most likely recovering gracefully within a few minutes.
    - Fault Tolerance in VMWare's VSphere is not a viable option at this time since that feature supports only a single virtual CPU. It is not feasible to expect either Prinergy or InSite to deliver required performance on a single virtual CPU system.
- **Network redundancy** from an Internet Service Provider may involve implementing multiple internet connections from multiple vendors to ensure uptime requirements. Since different vendors will have access to different public IP addresses and ranges, failing over in this scenario will require at least some manual configuration and possibly custom development. This area does not present anything unique for the virtualized environment that is different from a physical environment and so will not be discussed at length here.
- **Power Redundancy** is accomplished through implementing a range of possible strategies, discussed below. Again, there is nothing unique to the virtualized environment so it will be a limited discussion.

- Uninterruptable Power Supplies provide a battery backup for short-term to mid-term power failures. These might include both small system-based UPS for short-term (such as what might be installed locally in a rack) and building-based UPS cabinets (for mid-term coverage).
- Diesel Generators provide longer-term power in case of a localized failure for a longer period of time.
- Multiple substation power feeds may be advisable in some areas for protection against regionalized grid-level failures.

Potential Redundant Architectures

Redundancy within a Single Site

The most basic form of redundancy is to simply provide multiple hosts within the same site environment. This provides continuity in the case of a host hardware failure and nothing more. The following diagram shows before and after states with a host server that fails. In a live system, the active Virtual Machines would probably be balanced across both servers and each server would contain enough resources (CPU and memory) to accommodate the Virtual Machines from both servers if needed. This form of failover redundancy is managed by the VSphere environment and restarts the virtual machines on the new host once the original host fails. As a result, users do experience a short period of downtime during the boot-up cycles but any applications that can autostart should gracefully recover.

Figure 1 Host Failover using High Availability

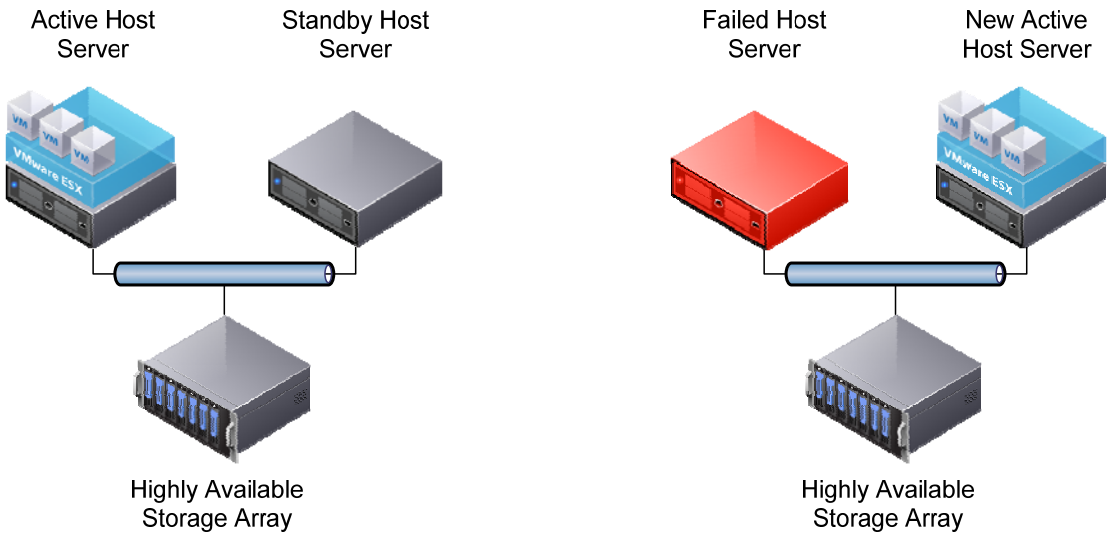

This failover can be set up within a single site or across two sites, but since there is a single storage location any site-wide failure at the storage location would lead to downtime.

Redundant hosts in multiple locations

A more effective form of redundancy is to have complete systems replicated in two locations, either within the same site or two different sites. This configuration requires that both the hosts and the storage be redundant,

and that the storage supports some form of replication. Many storage solutions support this form of replication either as a basic or a licensed feature.

Using this configuration, a system is protected against a host failure, a storage failure, or catastrophic site failure. VMWare's Site Recovery Manager can be used to script a managed failover from one site to another, or a manual procedure can be developed to point to the secondary storage location. There are several other factors, such as network addressing if the failover data center has public internet connections that must be configured, so this is not a solution out of the box but requires some planning and effort to define and enable if needed.

Figure 2 Muti-site redundancy with relicated storage

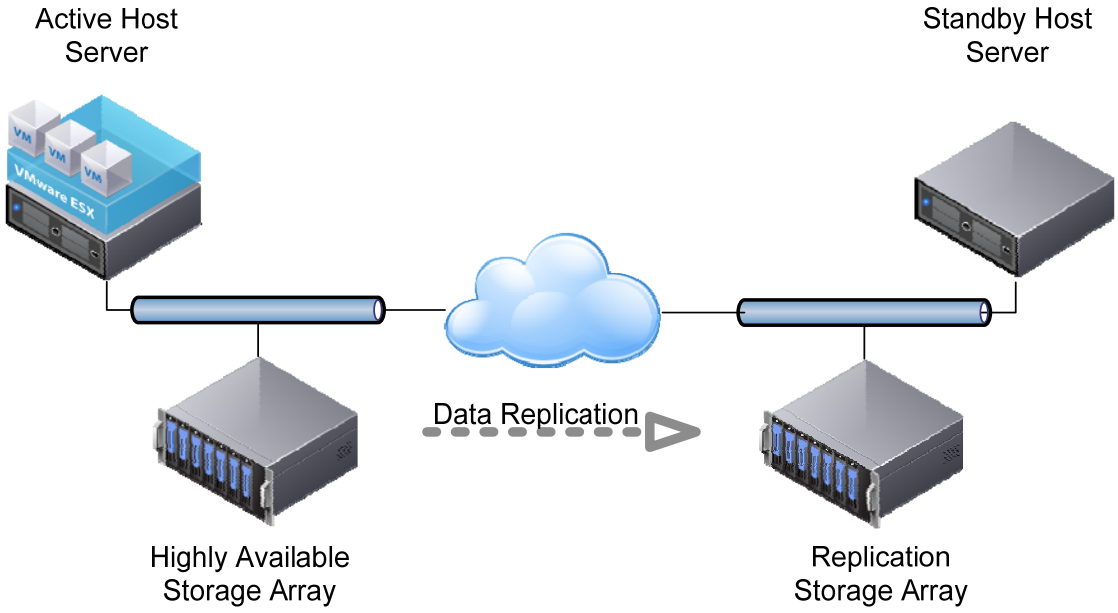

Virtualization is an important buliding block

It is important to recognize that the virtualized environment provides a framework over which managed or potentially automated fail-over solutions can be built, but these capabilities are not part of an out-of-the-box solution and they require very detailed and business-specific planning and work.

If you are interested in investigating a strategy and solution that meets your business continuity needs, Kodak can provide custom development resources to help you accomplish your needs. Please contact your sales representative for further details on Kodak's custom development opportunities.

About Kodak Graphic Communications Group (GCG)

The leading provider of graphic communications solutions worldwide, Kodak's Graphic Communications Group (GCG) offers image capture systems; professional color, copydot, and high-speed document scanning systems; inkjet printing and proofing systems; workflow and color management software; thermal imaging devices for film, plates, and proofs; high-quality proofing media, printing plates, and recording film; on-demand color and black-and-white printing systems; data storage products; and professional services. With corporate headquarters located in Rochester, N.Y., Kodak (NYSE: EK) has the largest global sales force and is committed to a digitally oriented growth strategy. For more information, visit www.graphics.kodak.com.