EXECUTIVE SUMMARY

Virtualization

DEFINING VIRTUALIZATION

A virtual infrastructure facilitates the full utilization of the physical resources of multiple servers across the infrastructure as a single unit, breaking the “one to one” relationship between applications and hardware. The virtual hosting model enables sharing the resources of a single physical computer across multiple virtual servers for maximum efficiency and reliability.

Individual applications can leverage as much or as little of the total computing capacity as demand dictates. A virtual environment can automatically move workloads from one workspace to another, prioritizing business needs while maximizing server resource utilization. By consolidating multiple physical servers to a single virtual environment, datacenter administrators can condense the number of servers they must support, reducing the overhead associated with disaster recovery, administration, maintenance, and power and cooling.

Virtualization improves the computing power/cost ratio, enabling a "pay-as-you-grow" scalable enterprise model to control capital expenditure and support business continuity objectives via the delivery of a highly elastic business solution.

TRADITIONAL ARCHITECTURE

Traditional system design directly links the operating system and application to a dedicated server and storage. This creates difficulty in scaling to meet growth, as the typical model for supporting increased demand is to purchase faster replacement hardware to replace the old equipment. For less active applications, this model also results in underutilized servers and storage.

VIRTUALIZED ARCHITECTURE

Virtualization breaks the one-to-one relationship between applications, hardware, and storage, allowing resources to be shared and allocated based on application demand. The virtualization control software functions as a traffic controller between the application and the hardware, continually evaluating application demand and allocating processing resources to meet the current demand.

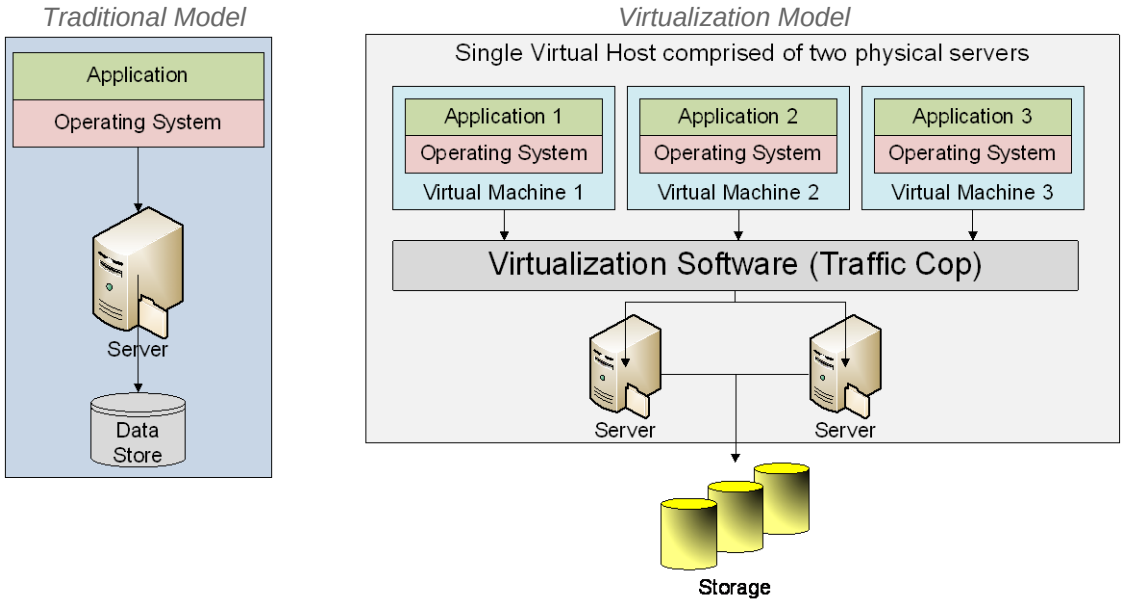

Figure 1 – Virtualization Architecture

Adding Computing Capacity

In the traditional, non-virtualized computing environment, when a server is no longer able to support the load generated by its application, a new and more powerful server is purchased and the information technology team migrates the application to the new hardware. The old server is subsequently decommissioned and removed from production. The new server must be capable of supporting the entire workload and is purchased with enough computing capacity to support future demand.

In the virtual model, a new server is added to the existing infrastructure to increase total computing power, with the old servers retained in production. Virtualization extends the lifecycle of server equipment by adding capacity to an existing server farm. This replaces the traditional model of replacing the underperforming server with a more powerful system before full depreciation and desired product lifecycle has occurred.

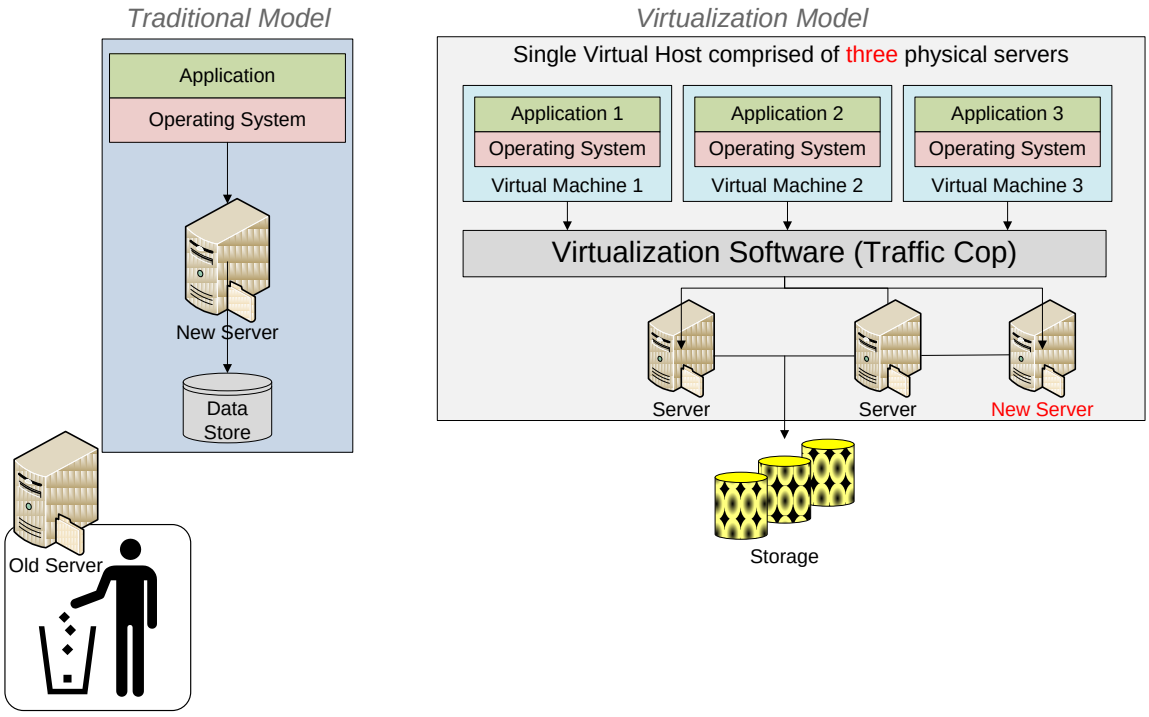

Figure 2 – Adding Capacity

High Availability Architecture

In a non-virtualized computing environment, when a server becomes unavailable due to an unplanned outage or planned maintenance, the application it supports becomes unavailable to the users for the duration of the outage.

In the virtual model, the loss of availability of any server triggers the virtualization software to direct the virtual server and its application(s) to run on the remaining, fully operational physical servers. The ability to utilize a dynamic workload management system dramatically reduces outage risk and associated lost productivity.

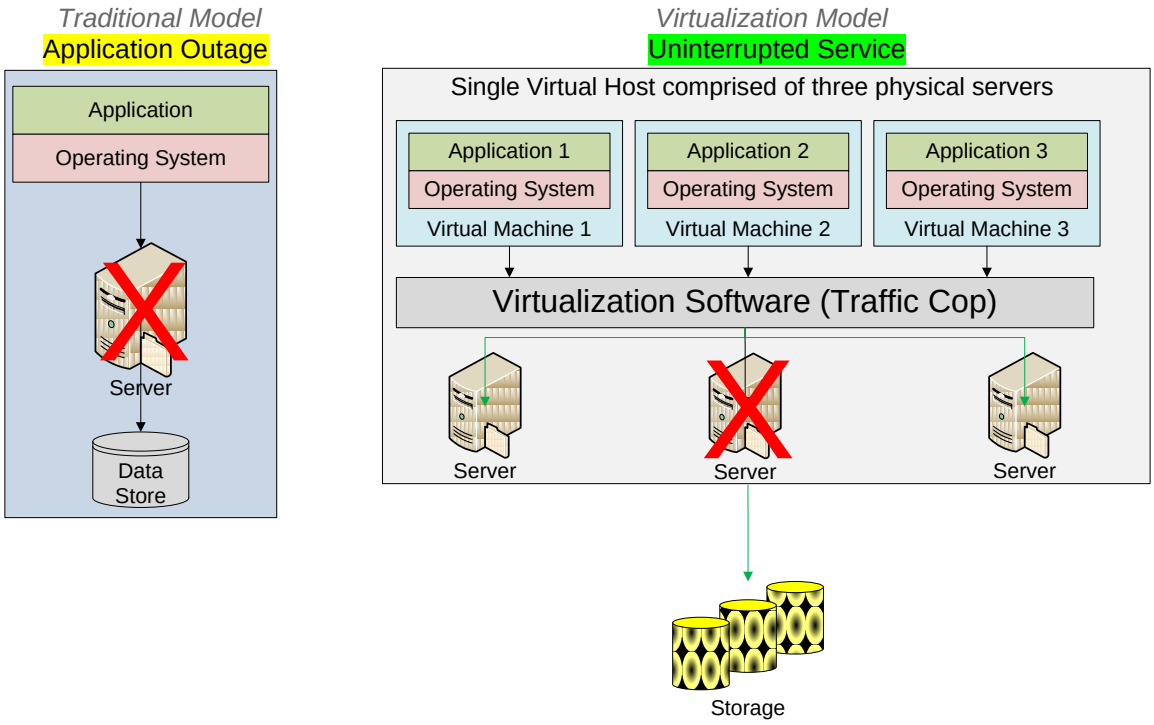

Figure 3 – High Availability

Disaster Recovery

The traditional Disaster Recovery model requires the information technology department to build, or in some cases, acquire and build replacement servers at the disaster recovery site, install and configure the operating system, and restore the application and related data from tape – a time consuming and high risk activity due to the difficulty in reinstalling software on diverse hardware.

The virtual model supports the utilization of moderately priced hardware to create a virtual server environment at the disaster recovery site. With the introduction of shared storage and replication technologies, the virtualization software can seamlessly transition the applications and data to the disaster recovery site with minimal interruption of service.

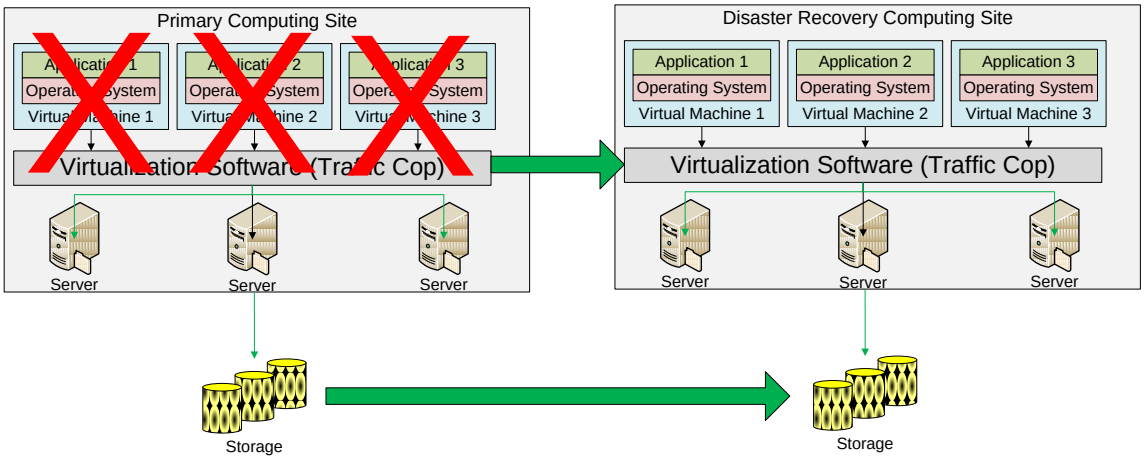

Figure 4 – Disaster Recovery

In practical terms, the introduction of virtual technology can reduce recovery time and recovery point objectives for key systems from hours/days to immediate/minutes, and the time to recover the secondary system from days/weeks to minutes/hours.

The following page highlights the dramatically simplified and accelerated recovery model available in a properly configured virtual environment.

Single Application Recovery Time
Physical vs. Virtual

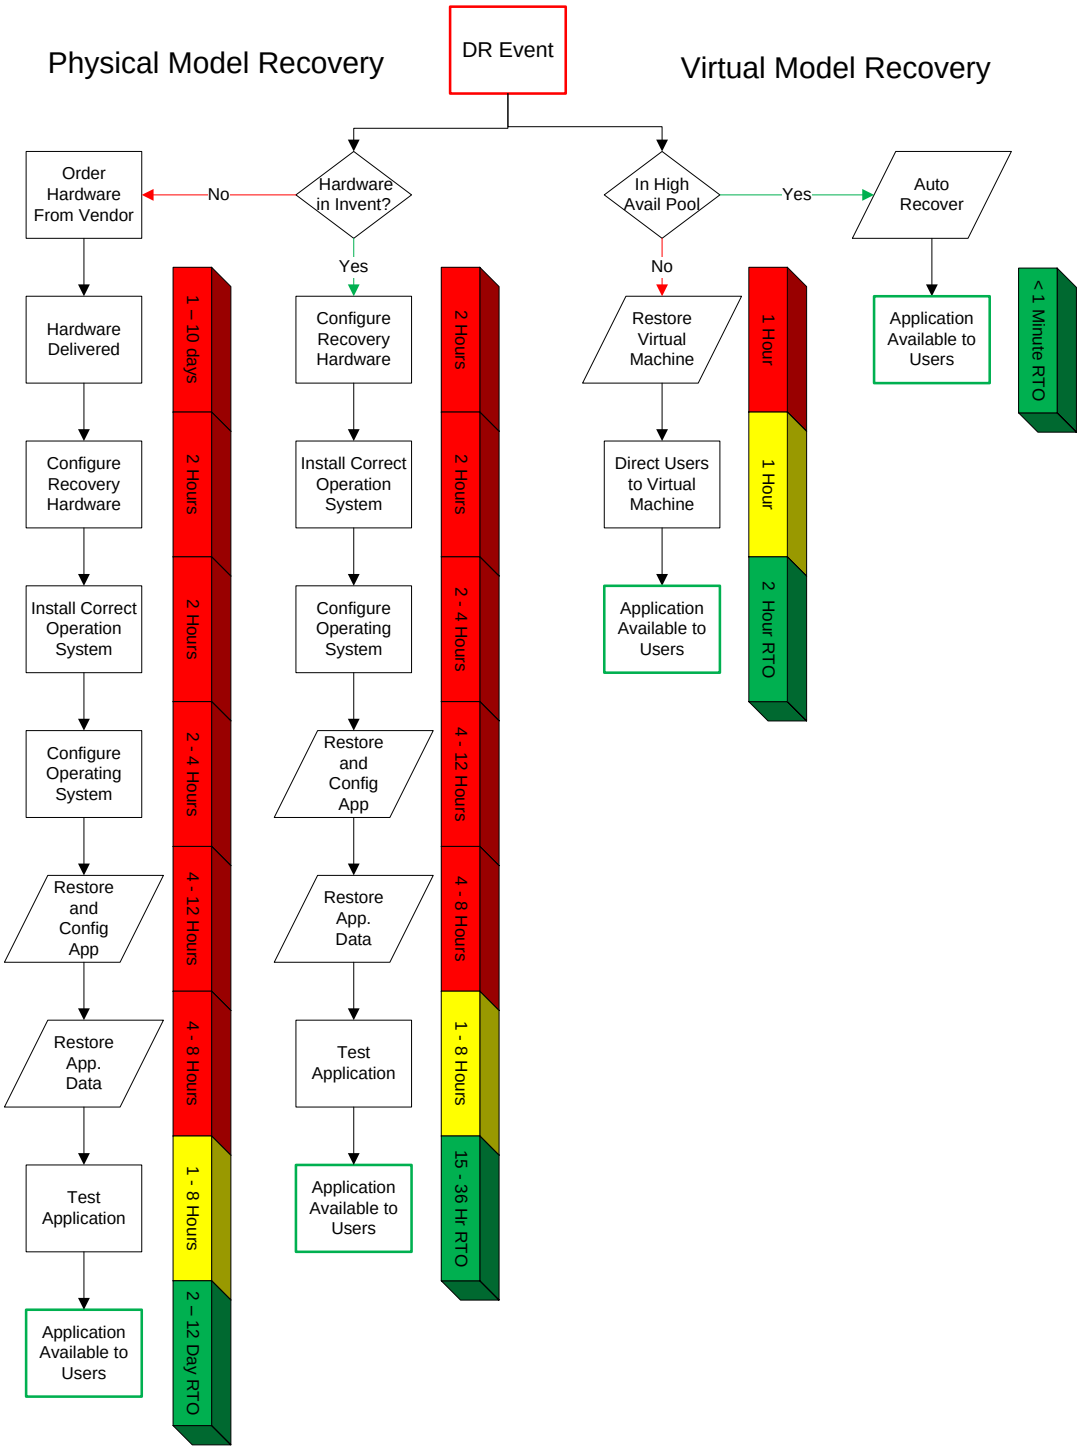

Figure 5 – Single Application Recovery Time Scenarios