

SERVER VIRTUALIZATION IN MANUFACTURING



Do's and Don'ts for Your Most Critical Manufacturing Systems



Abstract

While the benefits of server virtualization at the corporate data center are receiving a great deal of attention, its ability to address lifecycle management issues may make virtualization even more compelling at the manufacturing

plant. Along with the advantages, however, come additional challenges and risks. This paper offers best practices you can use to benefit from server virtualization today, while avoiding mistakes that could affect the availability and performance of business-critical manufacturing IT.

Server Virtualization For the Plant

Many companies are eager to jump into server virtualization, the practice of using a software layer to let one physical computing server run multiple virtual machines. Among many other advantages, server virtualization allows companies to save money by consolidating a number of applications on the same physical server.

Virtualization is a time-honored approach in the mainframe world, but its growing use today involves virtualization on x86 servers. The current surge of virtualization in the enterprise began with applications deemed less critical, characterized by lower processing requirements and tolerant of limited service outages.

Implementing virtualization in manufacturing environments moves the technology into the business-critical realm. Manufacturing Execution Systems (MES) and related automation applications can gain advantages from virtualization that go beyond even those seen in typical enterprise software applications — although special concerns apply as well.

Why Manufacturing is Defined as Business-Critical

Manufacturers are adopting new IT solutions on an unprecedented scale to meet efficiency, quality and regulatory compliance goals. And any time that you rely on complex integrated IT solutions in production environments, service interruptions quickly become unacceptable.

In applications such as MES, for example, unplanned downtime results in lost production time. In regulated industries like pharmaceuticals, loss of data and/or control can compromise the integrity of a batch record and require in-process product to be destroyed. Even minor system interruptions can call into question the value of the IT solution.

While application consolidation has many significant benefits for MES, significant risks exist when the underlying platform (including hardware, virtualization software layer and drivers) is not sufficiently robust. Platform faults could cause a facility-wide outage, and worse, long recovery time.

Using server clustering to provide a robust platform presents another challenge. Running a single application clustered on multiple servers in a non-virtualized environment is difficult enough. With virtualization, IT personnel have to deal with the complexity of configuring, testing and maintaining multiple applications that are clustered on the same platform.

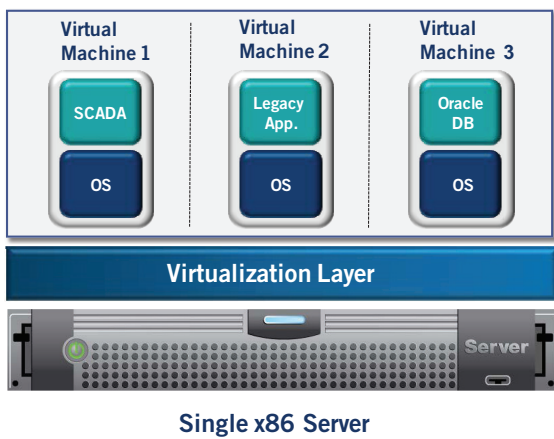
Therefore, evaluating and addressing an application's availability requirements is an important consideration before moving to a virtual environment. Before exploring these and other concerns in more depth, let's first review the basic concepts of server virtualization.

Server Virtualization Basics

In a virtualized environment, each virtual machine on a physical server exists within its own container or partition. While implementations differ, generally speaking each partition contains an application (or applications) and an instance of an operating system known as a guest operating system (OS).

A number of these partitions sit on a software layer called a hypervisor. The hypervisor is the thin, low-overhead layer that manages the basic services necessary to host the applications and their guest operating systems.

Figure 1: A View of Virtualization



Each virtual machine runs a separate instance of an operating system and application(s), and has access to a portion of the server's resources.

The virtual machines on a server may use the same flavor of operating system, use different releases of the OS or use entirely different types (e.g., Windows® and Linux®) of operating systems.

Some approaches use a host operating system below the hypervisor, but these impose system overhead. More recent solutions promote “bare metal” performance, which becomes possible when the hypervisor is implemented directly on the server hardware.

The market offers a number of commercially available software-based server virtualization products. The trend is to support virtualization that includes both Windows and Linux operating systems. VMware holds the largest market share. Other companies like Microsoft, Citrix and VirtualIron are strongly promoting themselves as contenders. In second place, Microsoft provides solutions for the Windows OS, and is working to extend support to the Linux OS. In addition, Microsoft and VMware have announced initiatives that would allow their virtualization solutions to work together, which holds promise for standardization.

Why Virtualize?

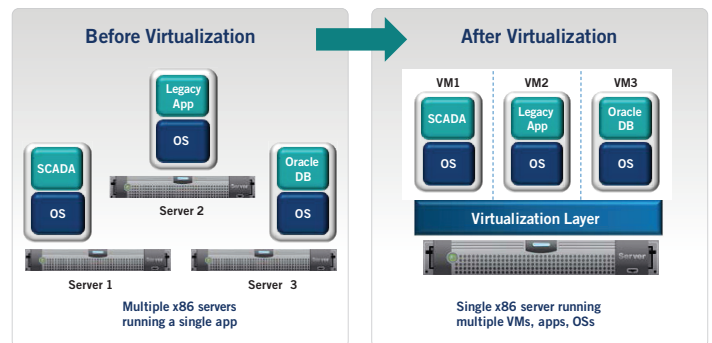
Reducing the number of physical servers tends to be the initial reason behind a decision to virtualize. Server consolidation is only the first of many potential advantages, especially considering the distinct characteristics of manufacturing IT solutions.

Server Consolidation

Many applications in the enterprise require as little as 5-10 percent of a server's capacity. Companies are opting to consolidate a number of these applications on a single physical server, while planning capacity for growth in the total workload.

At manufacturing plants, physical space tends to be especially tight. Facilities may or may not have computer rooms, so IT systems often have to compete for space with other kinds of equipment. Server consolidation can be especially beneficial at the plant level for these reasons.

Figure 2: Server Consolidation Example at the Plant Level



More efficient use of resources through server consolidation can bring savings in hardware-related costs and reduce environmental expenses as well: space, power and cooling.

Lifecycle Management

Compared with typical enterprise applications, MES and associated applications are distinguished by the need for long lifecycles. Lifecycles of seven years and up are common in the industry. After the IT solution is in production, companies want to ensure stability and reduce risk by avoiding changes to the software application, the operating system and the server hardware.

Achieving this objective becomes challenging because vendors often do not support the original operating system version throughout your desired lifecycle. This means you have to seek out extended support and pay a premium. Moreover, most server hardware is obsolete after three or four years.

Virtualization allows you to abstract the application and OS away from the server hardware. You can effectively extend the lifecycle of your application as a result. The ability of a hypervisor to support older guest operating systems allows you to upgrade the hardware platform without affecting applications or their operating systems. A related benefit is that the ability to upgrade server hardware eliminates the need to stock hard-to-obtain components required to maintain older computer servers.

Capabilities that virtualization can enable over the extended lifecycles of MES and related applications include:

- **Speed and ease of provisioning.** Server virtualization allows you to create a standard virtual machine — consisting of software files that include the application and an operating system — that can be copied onto a server in a matter of minutes when additional capacity is required, or when you need to distribute an identical application configuration to different plants. The virtual machine can be qualified and tested in advance to ensure it will work as expected. Besides the obvious implications for system stability in regulated industries such as life sciences, using a pre-validated virtual machine may eliminate platform qualification testing that would be required to install and validate new server hardware.
- **Hardware and capacity upgrades.** When more processing power or storage capacity is needed, virtualization can similarly let you move the virtual machine to newer hardware with no change to the application or operating system. You are not locked into older hardware.
- **Failover and disaster recovery.** A virtual machine's image — including configuration state, disk state and so on — residing on one physical server can also be periodically replicated to another physical server for backup or fast restart. Some virtualization software also allows for point-in-time rollbacks. Useful when data corruption has occurred, rollback lets an administrator revert the virtual machine to an earlier known good state.
- **Application development.** Taking as little as minutes to deploy, virtual machines effectively isolate each application developer in his or her own partition. Other developers are unaffected if a developer crashes the test application or guest operating system.
- **Upgrades without downtime.** A capability known as live migration allows for planned hardware and operating system upgrades (in cases where the operating system is not visible to the application) with virtually no interruption to the application and little perceived impact by users. Note that the operating system that can be upgraded is at the host OS/hypervisor layer; guest operating systems cannot be upgraded online.

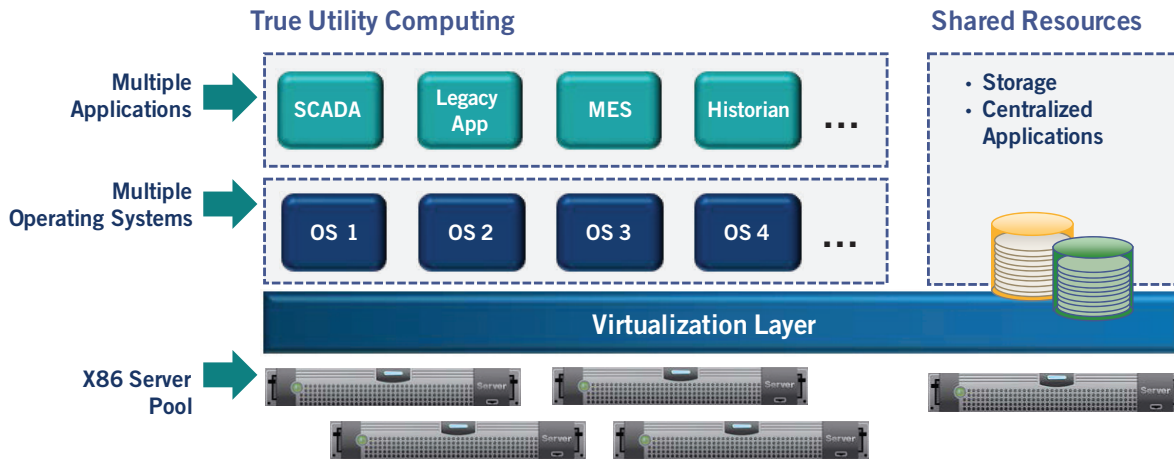
Live migration works by copying the system state iteratively while the application continues to run. Shortly before a final copy of the virtual machine is ready for migration, only a brief application blackout (perhaps milliseconds) is necessary to synchronize the second virtual machine with the original.

Ahead: True Utility Computing

As virtualization technology advances, users as well as vendors envision utility computing as the goal. People want pools of servers that not only run multiple applications, but also enable dynamic and fluid use of resources. In the future, automated management capabilities — such as application-sensitive monitoring and intelligent policy making — will be necessary to bring this flexibility to the mainstream.

Such utility computing will feature the self-management and policy making necessary to reduce requirements for IT labor and expertise. In addition, it will reduce the chance of operator error that industry experts cite as a leading cause of downtime.

Figure 3: Utility Computing Scenario



True utility computing enables resource sharing and dynamic resource allocation based on user-specified criteria, such as time of day and incremental capacity needs.

Figure 4: From Server Virtualization to Utility Computing

TODAY	TOMORROW
<ul style="list-style-type: none"> • Server virtualization is the focus • Requirement for IT expertise • Management and monitoring still developing <ul style="list-style-type: none"> - Ongoing IT administration is required - Policy making is not automated • Vendors have not implemented virtualization standard • Exposure to single point of failure at virtualization code layer and the server hardware 	<ul style="list-style-type: none"> • Utility computing is the focus • Simpler deployment, use will make benefits more broadly accessible • Management and monitoring come of age <ul style="list-style-type: none"> - Self-management - Policy making can be automated • Virtualization standards will increase ease of deployment, reduce risk • Single point of failure at virtualization layer is overcome by deploying in a fault-tolerant environment • Robust fault-tolerant hardware eliminates single points of failure on the virtualization server platform

Today taking advantage of everything that virtualization has to offer still requires the knowledge of, and management by, skilled IT professionals. And “putting all your eggs in one basket” can be a significant concern unless the underlying platform technology has the necessary robustness.

Best Practices for Your Business-Critical Applications

The risk and cost of service interruptions become higher as manufacturing applications become more powerful, integrated and relied upon by larger groups of workers. Employing best practices can help you achieve the advantages you seek from virtualization today, without compromising the availability and performance of your manufacturing application.

Know Your Application

Begin by characterizing your software application and its workload correctly. Which resources does your application consume? How much? When? How much headroom do you need for peak times and temporary surges in demand? In the event of performance degradation, the application could become unavailable and provide poor response time to users or processes.

Also make sure to conduct an appropriate risk assessment. Even if you are starting with non-critical applications, the server on which you are consolidating them often becomes essential when it drives numerous applications. In addition, not every application is a good candidate for virtualization. Typical examples are I/O-heavy applications and performance-sensitive environments that are not easily characterized.

Understand Tradeoffs

Because virtual servers are easy to set up and don't require the same management approval as hardware purchases, trade press articles report that some companies are experiencing "virtual server sprawl."

Expect some performance penalty as well; how much depends on your application and the virtualization technology you use. What's more, maximizing application availability and performance on a virtual machine requires considerable skill. And you need to understand how virtualization will affect your software license fees.

Seek Enterprise-Strength Technology

Remember that the virtualization layer has the potential to be a single point of failure for all of the virtual machines it supports. One rule of thumb: Software reliability increases as the amount of code and its complexity decrease.

Look for virtualization software that is small, compact, and controlled — as appliance-like in nature as possible. Virtualization and availability solutions that are simple to configure and maintain provide crucial advantages not only by reducing operating cost, but also by significantly reducing your exposure to downtime as a result of operational errors.

Plan for Business Continuity

Reliable availability and performance become more important the more that you depend on an IT resource, and the more that resource is integrated with other systems. To mitigate the risk of plant operations being interrupted, institute backup and disaster recovery measures for the physical servers that run your virtual machines.

Simplify With Robust Hardware

Virtualization subtracts physical complexity but adds equally real complexity in a virtual dimension. Without proper planning, this can be an issue because IT skills are in short supply at the average manufacturing facility.

Clustering multiple servers is one technique for achieving high availability. Implementing virtualization on a server cluster adds another layer to deploying and administering a cluster, on top of the ongoing attention from IT staff that is already necessary. For instance, a server must be running to migrate its workload to another member of the cluster. Therefore when a double-bit memory fault causes a server to crash, its workload cannot be transferred, data will probably be lost and a reboot will be necessary. And with clustering, you can expect to incur performance overhead in any case.

Server hardware or virtualization software vendors may claim they achieve high availability by predicting hardware faults and by enabling live migration of applications to backup servers; they may also assert that applications can be quickly restarted on another server. Ask which hardware faults can be predicted enough in advance to support a live migration; what percentage of the hardware is covered? In the case of a full restart, what is the worst-case fault detection and restart time? Is the management software making fail-over decisions robust, or for that matter, running on a robust platform?

For simplicity's sake, consider a fault-tolerant server that automatically protects reliability and availability without requiring changes to your business-critical application. This approach uses redundant components while appearing as a single server to virtualization and application software. Ideally, the emphasis should be on preventing downtime and data loss instead of simply on quick recovery.

Don't Let I/O Sink the Ship

Incompatibilities related to I/O interfaces are a known cause of system instability and performance problems. Establish that I/O devices and drivers are compatible with the virtualization technology you plan to use. Be ready, willing and able to resolve incompatibilities up front if you need to use legacy or proprietary I/O cards to access specialized plant equipment networks, as is common with supervisory control and data acquisition (SCADA).

Don't Go It Alone

Virtualizing in a manner that promotes the availability and performance of business-critical applications requires considerable expertise. Today, this goal is not realistic without a knowledgeable staff and/or a trusted professional services provider.

It is easy to make mistakes when doling out resources to virtual servers, for example. While allocating disk and memory among virtual machines is rather straightforward and fine-grained, CPU and network resources are another story. An incorrect assumption will quickly drag down a CPU- or network-intensive environment.

Conclusion

As server virtualization technology matures, it is becoming suitable for the much more exacting demands of business-critical manufacturing applications. Server virtualization can be a boon for managing the lifecycles of the many applications that make up an integrated MES environment and other critical manufacturing applications— with a few cautions. You can gain new capabilities and reduce costs, as long as you choose appropriate technology and plan properly.

About Stratus

Stratus Technologies is the leading provider of infrastructure-based solutions that keep your applications running continuously in today's always-on world.

Stratus always-on solutions can be rapidly deployed without changes to your applications. Our platform solutions provide end-to-end operational support with integrated hardware, software and services. Our software solutions are designed to provide always-on capabilities to applications running in your chosen environment – physical, virtualized or cloud. Our approach and our people enable us to identify problems that others miss and prevent application downtime before it occurs. Multiple layers of proactive diagnostic, monitoring and self-correcting services are backed by a global team of engineers who provide immediate support no matter where in the world your system is located.

If **always-on** is an application requirement, Stratus Technologies has a solution that fits.