



# Simplifying Downtime Prevention for Industrial Plants

A Guide to the Five Most Common **Deployment Approaches**



# Simplifying Downtime Prevention for Industrial Plants:

A Guide to the Five Most Common **Deployment Approaches**

## Introduction

Preventing unplanned downtime is a growing concern in today's always-on world, and it raises a key question: You want to keep critical applications up and running, but how do you know which availability solution is right for your organization? In our [Downtime Prevention Buyer's Guide](#), we discussed the six key questions that can help you isolate the right availability protection for your applications. In this guide, we'll discuss the pros and cons of different approaches to preventing downtime to help you determine which option would work best for your organization.

Dozens of technologies are available for maintaining application availability, ranging from backup and restoration to data replication, application failover, and restart. Despite their differences, though, they are all reactive solutions. In other words, they can't prevent an outage, which means the loss of data is a very real risk and ongoing processing may come to a stop. Ultimately, you can measure the value of these solutions by comparing how little data is lost (the Recovery Point Objective, or RPO), and how quickly the affected application(s) comes back online (the Recovery Time Objective, or RTO).

One overlooked alternative in maximizing RPOs and RTOs is a continuously available solution. Historically classified as fault tolerance technologies, these ultra-high availability offerings are the only solutions capable of delivering combined RPOs and RTOs of zero. Organizations will frequently disregard these solutions due to the misconception that they are proprietary in nature, impose a crippling amount of performance overhead, or are prohibitively expensive.

## Solution Options

Avoiding unplanned downtime is an imperative in most industrial plants. However, several factors contribute to a solution's overall suitability for specific use cases. Key criteria to evaluate include:



Application availability



Loss of data



Complexity/ease of use



Fit within existing environment



Recovery time



Lifetime value/Return on Assets/Total Cost of Ownership

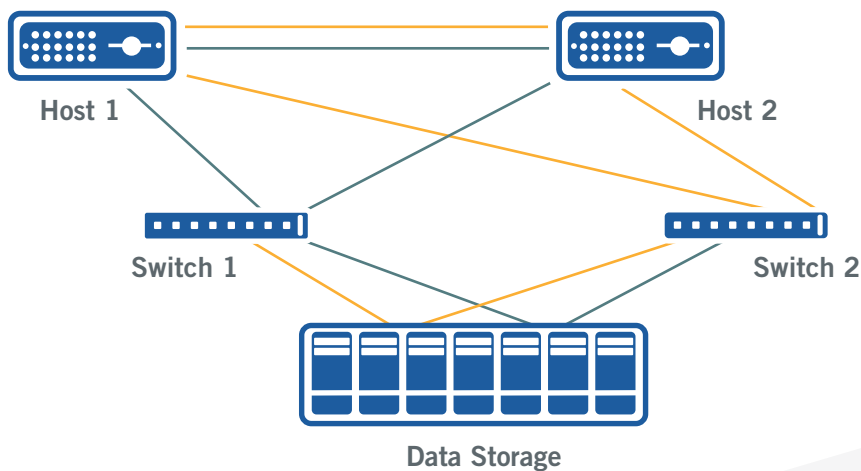
The [Stratus Downtime Prevention Decision Guide](#) describes each of these factors in detail.

This document discusses the advantages and disadvantages of each solution, using key selection criteria that will help determine which options meet the unique needs of your organization. We did not include standalone servers as part of the comparison because they can't deliver the availability service levels required for critical plant operations.

## OPTION 1:

### High-availability clusters such as Microsoft Windows Failover Clustering

Failover Clustering is a feature of the Microsoft Windows operating system that enables the grouping of servers to increase the overall availability of the application environment. While the nodes that make up the cluster are independent servers, they are connected through a combination of hardware and software. If a node incurs a failure, its workload is either restarted or failed-over to the remaining servers in the cluster.



#### Pros

- › Delivers high levels of availability if properly configured and managed.
- › Offers increased protection for virtualized environments.
- › Supports geographically dispersed multi-site configurations to provide disaster recovery capability.
- › Provides large number of nodes (sixty-four per cluster) with rolling upgrades/availability during planned downtime.


#### Cons

- › Mandates a reactive approach to availability that cannot prevent an outage from occurring.
- › All in-flight data is lost when an outage occurs.
- › Recovery times can range from minutes to hours based on factors such as affected hardware components, failure detection time, database recovery time, and application restart time.
- › A complex solution that depends on skilled personnel who must properly plan initial installation and enforce disciplined change management.

## OPTION 2 :

# High-availability solutions such as VMware HA for virtualized environments

VMware vSphere® High Availability (HA) is a feature bundled into most versions of VMware's vSphere® offering. Like Microsoft's Failover Clustering feature, HA enables a group of servers to offer a higher level of availability than standalone servers. vSphere HA clusters are composed of virtual machines (VMs) and their host servers, and they use VMware vCenter Server® application to monitor the state of the protected virtual machines. When a host failure occurs, the VMs on the affected server are restarted on surviving cluster nodes.



## Pros

- › Responds automatically to server failures, and quickly restarts VMs on surviving cluster nodes.
- › Restarts individual virtual machines and specific applications (when configured properly).
- › Implementation does not require application modification or failover scripts.
- › Included as a feature with most versions of VMware vSphere®.


## Cons

- › Mandates a reactive approach to availability that cannot prevent an outage from occurring.
- › All in-flight data is lost when an outage occurs.
- › While VMware vSphere® High Availability can respond quickly to a failure, actual application recovery time depends on the total time of the dependent services, such as the OS boot time and database start and recovery times.
- › Planning is required to ensure sufficient resources are available to restart virtual machines on surviving nodes following a failure.
- › The HA cluster relies on a shared storage device, which increases costs and administrative complexity.

## OPTION 3 :

### Fault tolerant solutions such as VMware vSphere® Fault Tolerance for virtualized environments

VMware vSphere® Fault Tolerance (FT) is a VMware solution that provides a higher level of data protection than their HA offering. FT delivers continuous availability by creating duplicate virtual machines hosted on separate physical host servers. Because the secondary VM's execution is identical to the primary, a failure on the server that is hosting the primary VM incurs no downtime. Failover is transparent, processing continues, and no data is lost.



#### Pros

- › Responds to server failures with zero downtime and no data loss.
- › No application modification required.
- › Fast setup.
- › Can be deployed with VMware vSAN™ software-defined storage solution.


#### Cons

- › Strict limitations on the maximum number of fault tolerant VMs allowed per host (four) and the maximum number of vCPUs aggregated across all fault tolerant VMs on the host (eight).
- › Expensive to configure for the maximum number of vCPUs (four per VM). Customers must license VMware vSphere® Enterprise Plus Edition™. Otherwise, only a maximum of two vCPUs per VM are supported.
- › Workloads that are latency-sensitive are poor candidates for deployment with VMware vSphere® Fault Tolerance.
- › While it will work with two nodes, VMware recommends that a cluster should consist of three hosts. All hosts must be licensed with the appropriate copy of VMware vSphere®.

## OPTION 4 :

### Stratus ftServer® fault tolerant solution

The Stratus ftServer® is a family of servers based on industry-standard hardware components that host Microsoft Red Hat Linux and VMware vSphere® environments. What distinguishes the Stratus offering is its ability to provide a continuously available processing environment through the use of a dual-modular redundant architecture and lockstep technology that replicates all computational activity on two CPU-memory motherboards. In addition to providing uninterrupted application processing and protection of in-flight data, the Stratus offering insures all I/O activity using replicated PCI buses, adapters, and storage devices.



Fault Tolerant Host

#### Pros

- › A hardware fault-tolerant solution specifically engineered for continuous availability and operational simplicity across bare-metal and virtualized environments.
- › No application modification required; minimal performance overhead incurred.
- › No strict limitation on the number of virtual machines or vCPUs supported per server.
- › A provider with a proven corporate focus and thirty-seven years of experience delivering availability solutions.

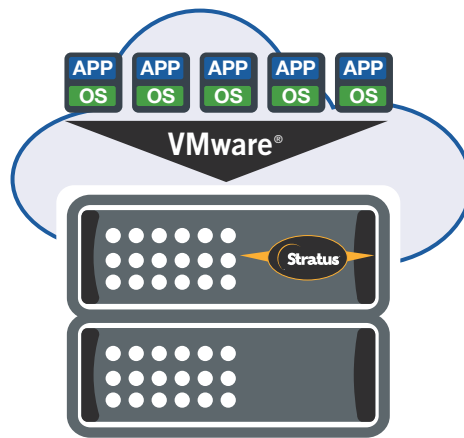
#### Cons

- › Limited product line (only three server models).
- › Higher initial acquisition cost compared to some alternatives.
- › The provider's niche status may create issues related to vendor approval status for some customers.
- › The provider has a smaller corporate profile than major competitors such as Dell™, HPE™, and Lenovo®.

## OPTION 5 :

# VMware + Stratus ftServer® fault tolerant solution

The Stratus ftServer® offers a unique yet complementary approach to providing an ultra-high level of availability for a VMware vSphere® environment. Unlike alternatives such as VMware's HA and FT offerings that duplicate infrastructure by requiring multiple servers configured with multiple copies of the software and hypervisors, the ftServer® implements a redundant architecture within a single standalone server, thereby reducing management complexity and licensing costs. Ultimately, the combination of VMware hosted on the Stratus ftServer® offers customers the optimum environment for deploying mission-critical virtual machines.



Virtualized Fault Tolerant Host

## Pros

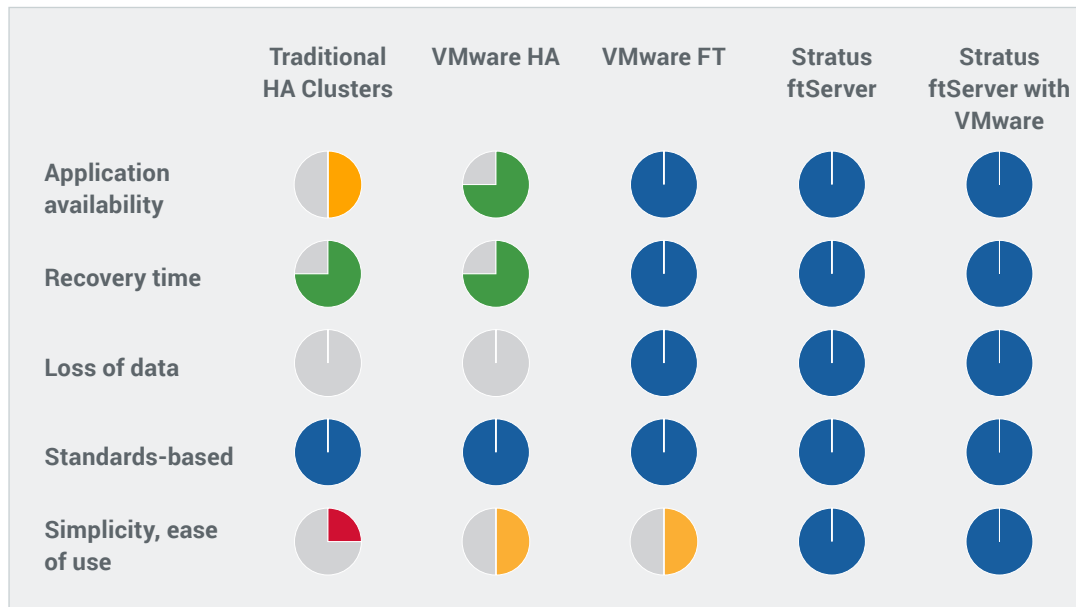
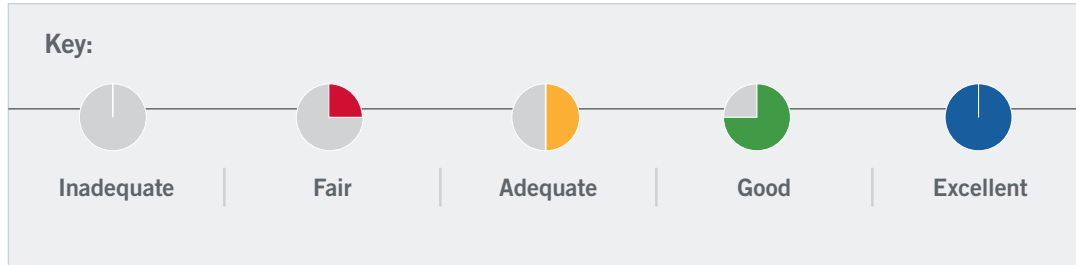
- › The combination of Stratus® and VMware® delivers superior virtual machine availability.
- › The single-server footprint simplifies administration and ongoing management, while deployment on a single server reduces software licensing fees.
- › Superior performance versus VMware vSphere® Fault Tolerance; no hard limitations on the number of fault-tolerant VMs/vCPUs per VM supported.
- › A provider with a proven corporate focus on availability solutions since 1980; providing continuous availability to VMware vSphere® environments since 2008.

## Cons

- › A limited product line (only three server models).
- › Some scalability issues at the high end.
- › May incur higher initial acquisition costs compared to some alternatives.
- › The provider's niche status may create issues related to vendor approval status for some customers.

## Availability Solution Comparison Overview

Each deployment option has been rated based on key solution criteria, including availability, recovery, loss of data, simplicity, and more. The chart below shows the relative performance in each criterion for all of the downtime prevention options listed above.





## The Stratus ftServer® Difference

Nothing else delivers what Stratus® can. Stratus' solutions far surpass other options due to their simplicity, ease of use, availability, and ability to integrate seamlessly into your environment.

Stratus offers both hardware and software solutions that meet the definition of a continuously available product. Both of Stratus' continuously available solutions are capable of delivering an RPO of zero. In other words, the applications running on these solutions incur no interruption or losses of data, even with a catastrophic event.

Stratus' solutions eliminate the loss of in-flight data through the use of redundant CPU/memory units. All instructions are executed in lockstep on duplicate components, so in-flight data will continue to process even if the worst occurs. And Stratus hardware and software solutions run industry-standard databases and applications without change or modification — so they integrate seamlessly into your existing computing environments. Stratus applications do not require modification and can be deployed out-of-the-box. Systems can be installed and managed without specialized training, allowing users to leverage the skillset of their existing IT staff. The results are simplicity of deployment and ongoing management for your organization.

The duplicate hardware and advanced software technology employed in our solutions allow the installation to remain “as-is” for a period that can stretch for years, regularly exceeding the industry norm. In addition to the obvious financial benefits that come with an application that is always-on, the avoidance of the capital investment associated with a hardware refresh, re-licensing of software, and the reduced administrative workload of the technical staff combine to deliver a low total cost of ownership.

### Take the Next Step

Any of these options can improve application availability. The key is to understand the relative importance of each of the six key criteria for your organization. Stratus can help you with this evaluation.

**Contact us for a free consultation on your availability readiness.**

Corporate  
Tel: 800.787.2887  
Website: [www.stratus.com](http://www.stratus.com)

Americas  
Greg Paden  
Tel: 978.461.7579  
Email: [greg.paden@stratus.com](mailto:greg.paden@stratus.com)

