

Stratus® everRun®

Affordable Software for Continuous Availability




Stratus everRun is downtime prevention software that ensures business continuity and data integrity by keeping Windows® and Linux® applications up and running with mainframe like availability.

Unlike other solutions, everRun prevents downtime, rather than merely recovering from it. And that difference has a big impact on your revenue, cost, customer satisfaction goals and efficiency rates.

everRun is easy to use. With click and go installation and centralized management tools for complete command and control of your IT world, it couldn't be simpler. And it's comprehensive monitoring and notification features is just one reason we call it "worry-free computing".

With everRun, you can use industry standard x86 servers so you can utilize existing competencies within your organization, without requiring specialized IT skill sets. This availability solution can run both your single and multi-threaded Windows and Linux applications fault-tolerantly, with no changes to your applications.

In the data center or at the edge of the network, in any industry, in any size organization, everRun is the solution you can rely on to meet your ever increasing demands in this always-on world. Stratus is the proven leader in preventative application availability and has been for over 30 years.



Key Benefits of everRun


- Prevents application downtime locally and cross-campus
- Enables business continuity and compliance
- Ensures no business interruptions or in-flight data loss
- Provides integrated disaster recovery
- Works with off-the-shelf industry standard servers
- Installs easily and affordably without the need for specialized IT staff
- Requires no modifications to applications
- Operates in virtual and cloud environments
- Protects VMs, applications and data
- Adapts to changing business needs with a scalable system
- Provides fault-tolerant or high availability protection for symmetric multiprocessor (SMP) and multi-core workloads

everRun specifications

Server	A variety of industry standard x86 systems are supported.
RAM	Up to 384 GB
Virtual machines	Up to 28
Supported operating systems	
Microsoft®	A variety of Microsoft Windows Server and Desktop guest operating systems are supported.
Linux®	A variety of Red Hat, SUSE, Ubuntu, Oracle and CentOS Linux guest operating systems are supported.

How does everRun work?

One application lives on two virtual machines. If one machine fails, the application continues to run on the other machine with no interruptions or data loss. If a component fails, it's replaced with the healthy component from the second system. The I/O is automatically mirrored to the redundant server. Memory check-pointing ensures all in-flight transactions as well as data in memory and cache is preserved — without machine restarts.



Comparison of VMware vSphere® and Stratus everRun

Complexity vs. Simplicity in Implementation

VMware vSphere® HA

Establishing a high-availability environment based on clustering is often a complex and costly endeavor. VMware vSphere® HA requires, at minimum, a two-node cluster to enable failover protection for VMware environments (Figure 1). Currently, VMware FT creates a shadow instance of a VM, requiring multiple hosts and additional storage configuration. VMware FT also requires a dedicated network for VMware vMotion to enable failover.

In a VMware environment, each node belonging to the HA/FT cluster must be identified and appropriate network connectivity between them established. The solution also requires an external storage array. It is then necessary to configure a number of settings to determine how the HA or FT cluster will behave and ensure that adequate resources are available in the event of a failover.

(Source: VMware vSphere vs Stratus everRun Enterprise Comparison| White Paper)




Fig. 1

Stratus everRun®

Stratus everRun prevents downtime using two x86 servers and a software check-pointing architecture to create a continuous availability environment (Figure 2). The solution runs on Centos Linux® and uses the open source KVM hypervisor.

Availability Extensions are routines that work behind the scenes on behalf of virtual machines to perform I/O redirection and synchronization between the two servers, while access to the physical hardware, VM scheduling, and memory management use the Centos and KVM hypervisor layers.

everRun is installed on both servers, requiring few choices for the components since it automatically uses appropriate defaults. Application I/O is automatically mirrored to the redundant server, and memory check-pointing ensures that all in-flight transactions, as well as data in memory and cache, are preserved without requiring machine restarts.

Stratus everRun provides two selectable availability levels — FT and HA. FT protection features synchronized memory and equal application state from either host. There is no single point of failure or single point of repair in FT mode, and storage is fully mirrored between hosts with no need for external storage.




Fig. 2

Availability Report Card


In terms of maximizing availability of business-critical virtualized applications, vSphere HA and everRun offer two fundamentally different approaches: vSphere HA clusters rely on an application restart after a server failure while everRun provides downtime prevention. Because of these differences, there are several important factors to consider when evaluating the two solutions.

VMware vSphere

While vSphere HA can reduce application downtime by automatically restarting virtual machines upon system or VM failure, it does not protect the host server or hypervisor against downtime or performance degradation, and any data not committed to disk is lost when an outage occurs. Moreover, application restart is highly dependent on the type of application being run and may not be instantaneous.

Stratus everRun

With Stratus everRun, an application resides simultaneously on two virtual machines. If one VM fails, the application continues to run on the other VM with no interruptions or data loss. If a component on the host server fails, the application continues to run on the healthy component from the second server. Most important, all in-flight transactions and data are fully protected from loss.



Feature	vSphere HA	Stratus everRun
Unplanned Downtime	53 minutes	< 4 minutes per year
Unplanned Downtime Cost	\$601,155	\$45,750
Maintenance Downtime	Requires restart	None
Failure Recovery Time	Minutes	None
Data Loss	Yes	None
Recovery Effort	Minimal	Fully automatic; no restart required
Fault Detection	Yes, automatic	Yes, automatic

Choose the Version That's Right for You

	everRun Express	everRun Enterprise
Supported availability types	High-availability	High-availability and fault-tolerance
Checkpointing scope	Data written to disk	Data written to disk, in-memory data, CPU states
Availability level	99.99% (4 nines)	99.999% (5 nines)
Recovery time	Depends on application, typically under 10 minutes	Under 5 seconds

Key Features of everRun

Achieve always-on availability with the Availability Engine

The Availability Engine lies at the heart of everRun. Complete with selectable availability, you can designate fault-tolerant availability or high availability for each VM. everRun ensures you won't have business interruptions because your applications co-exist on two virtual machines, with no changes to your applications, data loss or machine restarts.

Know when to take action with the Watchdog Service

Stratus' support center constantly monitors your system and automatically sends you a system level notification should a fault occur.

Save time with The Stratus One View Console

Monitor and manage everRun instances across multiple deployments from a centralized location. Build, designate and deploy servers and Virtual Machines.

Simplify, accelerate and optimize performance management with Application Monitoring

Get complete performance visibility across dynamic IT environments. Know immediately when an application or device is down and have it restart instantly to minimize downtime.

Protect against localized disaster with SplitSite® Cross-Campus downtime prevention

SplitSite provides application fault tolerance across geographically separated sites using synchronous replication.

Mitigate disaster impact with Disaster Recovery

Get up and running faster than ever before with built-in asynchronous replication between sites over a wide area network connection.



ESTABLISHED IN 1987.

Servo Dynamics is a Singapore based enterprise, specialised in providing Industrial Automation and Control Solutions. Servo Dynamic's expertise covers a diverse range of products and services, including robotic systems, factory automation, drive and control technologies, mechanical components, industrial computing and safety systems.

Servo Dynamics core focus is to offer cost-effective, high-quality end-to-end solutions, emphasizing customer satisfaction, adhering to regulatory requirements and exceeding standards.



servodynamics@servo.com.sg

101 Defu Lane 10
Singapore 539222

www.servo.com.sg