

# Veritas InfoScale™ Availability

Delivers high availability and disaster recovery for your critical applications

## Overview

Veritas InfoScale™ Availability protects your most important applications from planned and unplanned downtime. InfoScale Availability provides high availability and disaster recovery for your critical business services including individual databases, custom applications, and complex multitiered applications across physical and virtual environments, and over any distance. InfoScale Availability uses intelligent monitoring for applications and infrastructure to detect possible risks to availability, and automatically recovers applications when needed. InfoScale Availability also detects site outages and initiates application recovery at a disaster recovery site. With built-in testing to proactively detect and report potential problems before they have an impact on IT services, InfoScale Availability helps minimize unwanted downtime of your critical business services.

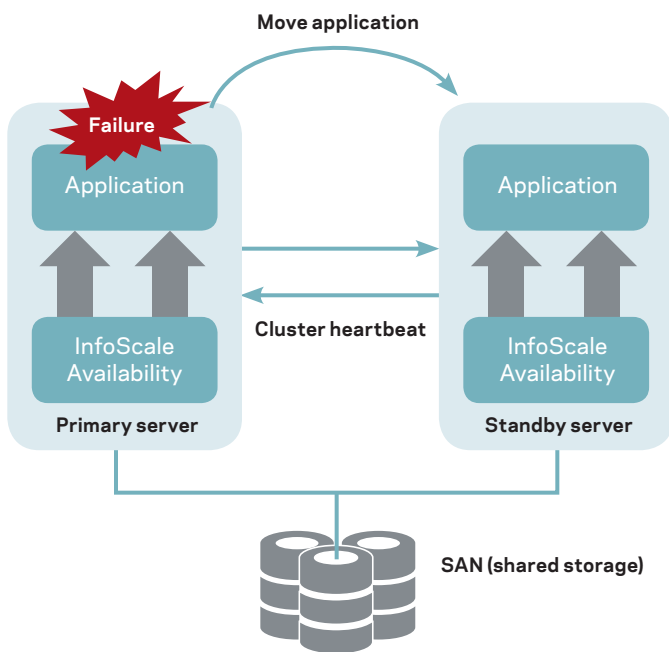


Figure 1. InfoScale Availability intelligently monitors applications and recovers automatically after a failure.

## Highlights

- **Veritas™ Virtual Business Services**—Provides automated and orchestrated recovery of applications composed of multiple components running on different physical and virtual tiers, adding resilience to complex business services. Virtual Business Services also supports Microsoft® Failover Cluster within a physical tier, as well as for Microsoft® Hyper-V guests in a virtual tier, giving faster and automated recovery for multi-tier applications with components across native Windows® clusters
- **Out-of-the-box support for applications and databases**—Guarantees compatibility for hundreds of applications and databases, reducing deployment times and consulting costs
- **Comprehensive hardware and platform support**—InfoScale Availability is the only solution that supports all leading operating systems, including Windows, Linux®, and virtual platforms including VMware® ESX®, Red Hat® Enterprise Virtualization (RHEV), Oracle® VM, and Hyper-V
- **Ensure application availability in VMware environments**—Failover applications without a VM reboot to standby healthy virtual machines, for fast application recovery from any failure without impacting VMware features such as VMware® vMotion™ and Distributed Resources Scheduler
- **Ensure high availability for Cisco® Unified Computing System Servers**—InfoScale Availability includes a high availability add-on agent for Cisco Unified Computing System servers, allowing service profiles to be restarted on standby-by blades in case of primary hardware failure
- **In metro VM availability for Hyper-V environments**—Failover mission-critical virtual machines running business applications across metro sites without any data loss (Zero RPO)

- **Faster failure detection with Intelligent Monitoring Framework**—Our Intelligent Monitoring Framework detects failures faster than traditional clustering solutions and requires almost no CPU overhead
- **Advanced failover logic with AdaptiveHA**—Ensures that resources are utilized as efficiently as possible by failing over applications or virtual machines to the most optimal server
- **Availability over any distance**—Builds both local and remote clusters for high availability and disaster recovery and works with all major data replication technologies
- **Automated disaster recovery testing**—Tests application failover without affecting the primary environment
- **Simple installation, configuration, upgrades, and management**—Provides wizard-driven installation for easier implementation and a central management console to view and manage multiple local and remote clusters. InfoScale Availability can be upgraded without impact to in production applications
- **Multiple cluster management and reporting**—Enables administrators to monitor, manage, and report on InfoScale Availability implementations on different platforms from a single Web-based console Veritas™ InfoScale Operations Manager
- **Data integrity protection**—Protects application data through advanced network arbitration and proactive prevention of application concurrency in failover environments

## Virtual Business Services

Today's data centers consist of multiple layers of physical and virtual environments, each with its own administrative tools and authorization requirements. This creates an end-to-end management challenge, with applications or entire business services composed of multiple components that interact with each other despite running on different operating systems and virtualization technologies. With the introduction of Virtual Business Services, IT administrators now have visibility across multitiered applications and can automatically recover them.

Recovering a single failed component of the business service does not guarantee business continuity. Not only does the failed component need to be recovered, but it must also be reintegrated into the complete business service to keep downtime to a minimum. Virtual Business Services, included with InfoScale

Availability and managed through InfoScale Operations Manager, is aware of the complete business service and can take action in the event of a failure. When an individual component of the service fails, InfoScale Availability will not only recover the failed application, but it will also automatically orchestrate the connection to other computing resources needed to keep the business service available. The end result is faster recovery and minimal downtime—all with no manual intervention.

Virtual Business Services also supports Microsoft Failover Cluster running on one or more business service component layers, providing increased availability control for critical business services across heterogeneous infrastructure and clustering solutions.



Figure 2. Virtual Business Services - the building blocks for your private cloud

## Out-of-the-box support for applications and databases

Installing other clustering solutions often requires long consulting projects and custom scripting to support the different applications and databases in your environment. InfoScale Availability provides off-the-shelf support for a wide range of applications, including SAP®, BEA®, Siebel®, Oracle applications, Microsoft® Exchange, and PeopleSoft®, as well as enterprise-class databases such as Oracle, DB2®, Microsoft SQL Server®, and Sybase®.

Using InfoScale Availability means less scripting, faster installation times, and easier ongoing maintenance when upgrading your applications. In addition, InfoScale Availability supports custom-built applications with generic or custom-developed agents.

### Comprehensive hardware and platform support

Multiple platforms often mean multiple high-availability tools, and with them, increased complexity. This complexity can increase administrative costs and present a higher likelihood of error. InfoScale Availability is the only solution that supports all leading operating systems, including Windows, Linux, and virtual platforms including VMware ESX, Red Hat Enterprise Virtualization, Oracle VM, and Hyper-V.

### Ensure application availability in VMware environments

InfoScale Availability enables IT organizations to monitor applications running inside Windows or Linux virtual machines residing on VMware ESX servers. InfoScale Availability monitors the application and, in the case of a failure, handles the recovery of the application to a healthy virtual machine running on a different ESX server without requiring a VM reboot. InfoScale Availability facilitates fast application recovery in response to a wide range of failures, while fully supporting the native VMware tools such as vMotion, Distributed Resource Scheduler and Site Recovery Manager. InfoScale Availability can be easily managed from a VMware vSphere® client through a VMware vCenter® plugin, and is easy to set up with its auto-discovery and auto-configuration capabilities.

InfoScale Availability provides an added layer of protection even when an operating system corruption issue occurs. In the instance of an operating system corruption issue, a restart of the application or virtual machine will not solve the problem. One option is to restore the last known good copy of a virtual machine, but this can be a time consuming and manual process. InfoScale Availability effectively solves this problem by providing the capability to recover the application in a new, healthy virtual machine on a different ESX server without a reboot. This results in much faster application recovery and less downtime. The same application failover capability can help minimize planned downtime during maintenance activities, such as OS patching.

### Ensure high availability for Cisco Unified Computing System Servers

InfoScale Availability has an agent that integrates with Unified Computing System Manager and allows automatic detection of service profile configuration and hardware related faults. The agent restarts a service profile on a failover target in case of a hardware failure with the original Server blade.

### Ensure application availability in Hyper-V, across the metro region

In Hyper-V environments, InfoScale Availability provides availability for VMs running critical business applications across metro sites, without any data loss (Zero Recovery Point Objective). This feature enables Hyper-V VMs to be configured in a single stretch cluster, with heterogeneous storage, and ensures there is no data loss when VMs failover across metro distances.

### Faster failure detection with Intelligent Monitoring Framework

Normal clustering relies on resource polling to determine the health of application resources. This polling process increases processor overhead but, more importantly, faults may not be noticed immediately. InfoScale Availability enables the faster detection of faults by asynchronously monitoring selected resources with its Intelligent Monitoring Framework. This means that failures can be detected instantaneously, instead of waiting for a nonresponse from a faulted resource. Furthermore, the CPU overhead associated with traditional poll-based monitoring is also reduced.

### Advanced failover logic with AdaptiveHA

Most clustering solutions recommend using active/passive two-node clustering configurations, which means that backup servers remain idle, wasting computing resources and decreasing server utilization. On the other hand, InfoScale Availability allows true N+1 “roaming spare” or N+M “active/active” capability for maximum availability without the cost of a dedicated spare per application.

For physical Linux clusters, InfoScale Availability ensures maximum uptime and optimal server utilization by using dynamic target system selection for application failover by determining the cluster node with the most available CPU and memory. It monitors the available capacity of systems in the cluster in terms of CPU, memory, and swap to select the most optimal target based on application needs. When a failure occurs, InfoScale Availability can automatically choose the least utilized server to failover to, as well as add repaired servers back into the selection pool when they rejoin the cluster.

For virtual environments (LDMs), InfoScale Availability monitors and forecasts the available capacity of the target physical host in terms of CPU and memory as well as the CPU and memory

capacity assigned to the virtual machine. AdaptiveHA enables InfoScale Availability to dynamically select the most optimal target physical host to online, switch and failover the virtual machine service group to.

### Availability over any distance

Disaster recovery is more than replicating data. For mission-critical applications that must remain online even in the event of a site failure, disaster recovery solutions should also automate application recovery. InfoScale Availability provides disaster recovery over any distance, allowing organizations to deploy both local high-availability and remote clustering for complete disaster recovery. With a single click, InfoScale Availability can migrate applications between servers in a local data center or move all applications to a data center thousands of miles away.

For campus clusters, InfoScale Availability provides intelligent failover over all associated service group dependencies for an application during a cross-site failover. For example, if an application is failed over, the supporting database is failed over as well to reduce latency and maintain performance speeds. Additionally, for stretch clusters, site-based priority can be set giving greater weight to the preferred site.

InfoScale Availability supports all major hardware, software, and database replication technologies. Veritas offers an integrated solution for application and data availability. Replicator Option technology, a feature of InfoScale Storage, provides continuous data replication that transfers data over any distance. In addition to Replication Option, InfoScale Availability provides full support for all major third-party data replication solutions, including Hitachi® TrueCopy, HP Continuous Access XP, HP Continuous Access EVA, EMC® SRDF, EMC® RecoverPoint, EMC® MirrorView, NetApp® SnapMirror, IBM® Metro Mirror, IBM® Global Mirror, IBM® HADR, IBM® XIV, Oracle Data Guard, and others.

Using InfoScale Availability for disaster recovery means faster recovery times, less reliance on personnel during a disaster, and simpler management of data and application recovery. InfoScale Availability completely automates the process of replication management and application startup at the remote site, without the need for complicated manual recovery procedures involving storage and application administrators.

InfoScale Availability provides all the necessary logic to completely control the underlying synchronous or asynchronous replication configuration.

### Automated disaster recovery testing

Production servers and applications are constantly changing, therefore, regularly testing a disaster recovery strategy is critical to guarantee a successful recovery in the event of an outage. To better enable successful recovery, InfoScale Availability includes Fire Drill, a tool that simulates disaster recovery tests by starting up an application at the disaster recovery site as it would in an actual disaster. Because it is a simulation, Fire Drill does not disrupt production applications, so it can be run as often as necessary, without the disruption and cost of traditional disaster recovery testing.

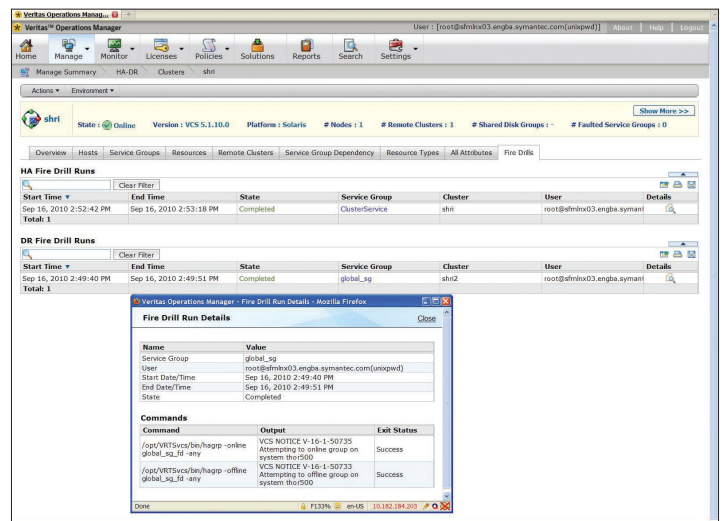


Figure 3. With InfoScale Availability, you can test your disaster recovery failover without disrupting production applications.

### Simple installation, configuration, upgrades and management

Deploying new clusters can be difficult with other clustering solutions, due to poor installation tools and heavy scripting requirements. InfoScale Availability provides administrators with easy-to-use configuration wizards for simplified cluster implementation.

Setting up applications is also quick with application wizards. InfoScale Availability automatically detects and configures an application running on a server through the InfoScale Operations Manager console. This makes the process of making applications highly available easier as application service groups, networking, and storage resources are included.

Administrators can perform an online upgrade of InfoScale Availability using the installer, while production applications are online. Applications can run seamlessly when the upgrade is in progress. Cluster modules will be stopped and failover capabilities will not be available during the upgrade process.

In addition, administrators can use the Services and Operations Readiness Tool, a Web-based tool that provides risk assessment services and automates many installation tasks. The Cluster Simulator, a free download, allows cluster administrators to simulate application failover scenarios and familiarize themselves with InfoScale Availability without risk. Cluster Simulator helps administrators simulate high availability environments from their laptops, where they can test multiple application failover scenarios without affecting production environments.

### Multiple cluster management and reporting

With the increasing number of applications and heterogeneous clustered servers distributed across multiple data centers, management of clusters can become complicated. Using InfoScale Operations Manager, customers can now monitor, manage, and report on InfoScale Availability implementations on different platforms from a single Web-based console. The InfoScale Operations Manager capabilities increase administrator efficiency by providing enhanced visualization of the managed clusters, centralized control for global applications, and complete reports of each application's availability status. It helps administrators avoid common cluster configuration mistakes, audit unexpected cluster configuration changes, and provides a standard way for administrators to detect and investigate cluster problems and track management history of all the managed clusters.

Through InfoScale Operations Manager, InfoScale Availability also features quick and easy configuration setup and maintenance of campus clusters or multi-site clusters, with enhanced visibility.

### Data integrity protection

When intra-cluster communication breaks down, it is possible that two systems in a cluster could try to write to the same storage and cause data corruption. The advanced data protection logic in InfoScale Availability shields data from becoming corrupted when a "split brain" situation arises by providing arbitration over cluster membership decisions. This guarantees data integrity as well as availability of service. The membership arbitration is provided using the SCSI III protocol utilizing an odd set of coordinator disks or, optionally, through a software solution that is enabled via a Coordination Point Server.

InfoScale Availability further prevents data corruption through the detection and proactive prevention of failover applications starting simultaneously on more than one server during a failover.

NOTE: All InfoScale Availability features are also available in InfoScale Enterprise.

## More Information

### Visit our website

<http://enterprise.symantec.com>

### To speak with a Product Specialist in the U.S.

Call toll-free 1 (800) 745 6054

### To speak with a Product Specialist outside the U.S.

For specific country offices and contact numbers, please visit our website.

## About Symantec

Symantec Corporation (NASDAQ: SYMC) is an information protection expert that helps people, businesses and governments seeking the freedom to unlock the opportunities technology brings - anytime, anywhere. Founded in April 1982, Symantec, a Fortune 500 company, operating one of the largest global data-intelligence networks, has provided leading security, backup and availability solutions for where vital information is stored, accessed and shared. The company's more than 20,000 employees reside in more than 50 countries. Ninety-nine percent of Fortune 500 companies are Symantec customers. In fiscal 2014, it recorded revenues of \$6.7 billion. To learn more go to [www.symantec.com](http://www.symantec.com) or connect with Symantec at: [go.symantec.com/socialmedia](http://go.symantec.com/socialmedia).

## Symantec World Headquarters

350 Ellis Street  
Mountain View, CA 94043 USA  
+1 (650) 527 8000  
1 (800) 721 3934  
[www.symantec.com](http://www.symantec.com)