aws marketplace

InfoScale

Reviews, tips, and advice from real users






Powered by PeerSpot

Contents

©2026 PeerSpot, All Rights Reserved

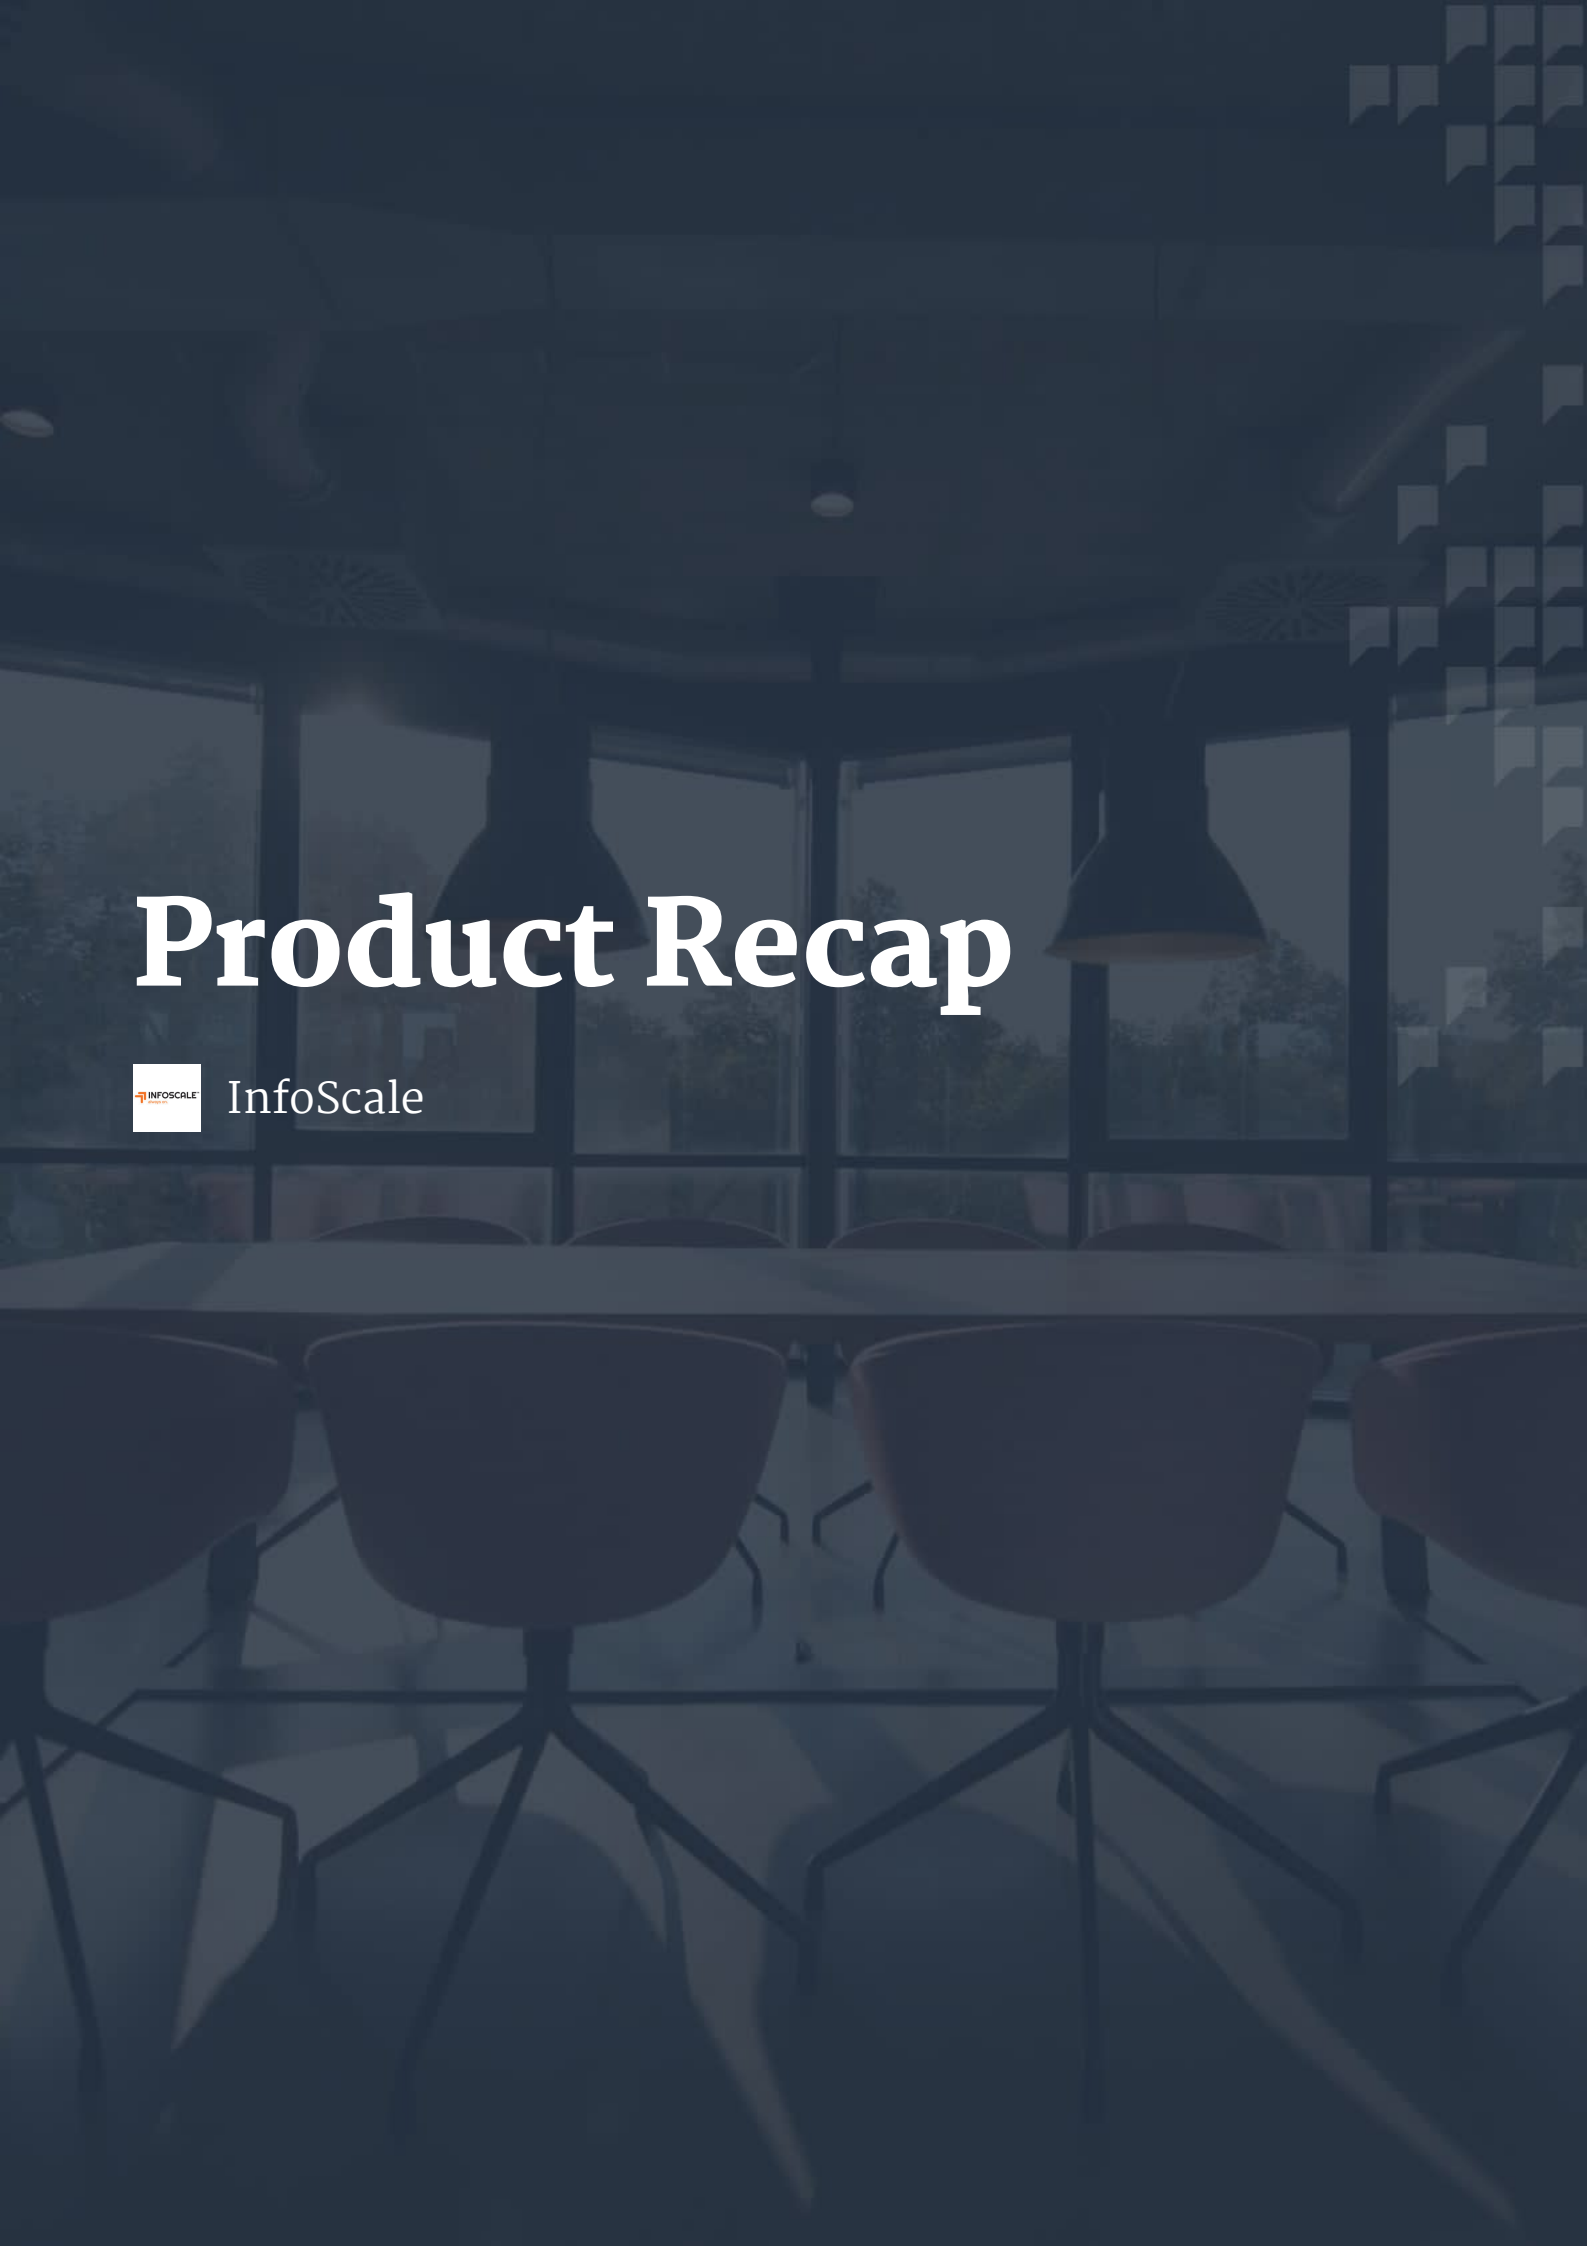

Product Recap

InfoScale

InfoScale Recap

InfoScale delivers robust data availability and protection for enterprises by managing critical workloads across hybrid cloud environments, ensuring minimal downtime and seamless performance.

InfoScale is designed to enhance data resilience and operational efficiency. By enabling quick data access across hybrid infrastructures, it reduces potential downtime and enhances business continuity. Its scalable architecture adapts to dynamic workloads, providing reliable performance under pressure while maintaining seamless integration with cloud-based solutions.

What are the key features of InfoScale?

- Data Resiliency: Ensures continuous data availability through advanced clustering and replication technologies.
- Scalability: Supports dynamic scaling to meet changing demands without compromising performance.
- Cloud Integration: Offers seamless transition and operation within cloud environments.
- Automated Operations: Streamlines management processes with automation to reduce manual intervention.

What benefits and ROI should users consider?

- Increased Uptime: Improved reliability translates to reduced business interruption.
- Cost Efficiency: Potentially lowers operational costs through optimized resource allocation.
- Enhanced Security: Provides robust protection measures for critical data assets.
- Operational Agility: Facilitates swift adaptation to evolving business requirements.

InfoScale is particularly effective in sectors such as healthcare and finance where data integrity and uptime are crucial. Its implementation empowers industries to handle data-intensive operations efficiently, ensuring continuity and compliance with industry standards.

©2026 PeerSpot, All Rights Reserved

Valuable Features

Excerpts from real customer reviews on PeerSpot:

- "From a recovery standpoint, InfoScale is excellent and easy to manage."

TABISH JAVED

Site Reliability Engineer (Certified) at Kyndryl India

- "InfoScale's ability to maintain data integrity and availability during a cyber event such as a ransomware attack is excellent."

Brad Waddell

Technical Trainer at netxinc

- "InfoScale does that all by itself without any dependency on different solutions."

SavasIrez

Professional Services Director – New Technologies at Gantek

©2026 PeerSpot, All Rights Reserved

"Over the last two years, we did not experience any application failover or receive alerts due to the immediate switchover mechanism in active-active mode that ensures no downtime, helping us significantly with confidence and trust in our organization."

Jaswanth Kotla

Sr Network and Security Engineer at a computer software company with 201-500 employees

"In a live incident scenario, the data replication process occurs in real-time, and compared to other products, this data replication feature works effectively, ensuring data availability, and we can implement this scenario using Veritas Volume Replication (VVR), which is the most usable feature in InfoScale for data replication."

Chathura Nuwan

System Engineer at MIT ESP

"The best feature is that it supports high availability and automatic failover, which minimizes downtime and helps the environment reduce downtime and improve high availability for critical applications."

Islam Hamada

Network Security Engineer at EMAK For Integrated solutions

©2026 PeerSpot, All Rights Reserved

"It offers High Availability for many applications, including Oracle and SAP environment."

Verified user

Storage and VMware Expert at a comms service provider with 1,001-5,000 employees

©2026 PeerSpot, All Rights Reserved

What users had to say about valuable features:

“The best feature is that it supports high availability and automatic failover. I can manage it in one way.

InfoScale automatically triggers failover, which minimizes downtime. Downtime and maximum uptime result in zero downtime.

It is good for high availability. It helps the environment reduce downtime and improve high availability.

It has reduced downtime and improved high availability.

Regarding resilience, I am the customer and we call the vendor to make the operation.

InfoScale provides high availability and automatic failover, and it improves the scalability for critical applications..”

Islam Hamada
Network Security Engineer at EMAK For Integrated solutions

Read full review

©2026 PeerSpot, All Rights Reserved

"The features or capabilities of InfoScale that I have found the most valuable and useful so far are high availability and service failover where we can create resource groups in the traditional cluster, which allow us to failover services and data volumes, and in InfoScale for Kubernetes, we can implement a DR solution that allows us to bring up the application from the DR side in case of a disaster at the production site.

In a live incident scenario, the data replication process occurs in real-time, and compared to other products, this data replication feature works effectively, ensuring data availability, and we can implement this scenario using Veritas Volume Replication (VVR), which is the most usable feature in InfoScale for data replication.

Data integrity in InfoScale is ensured by the VxFS journaling filesystem, application-level transactions, and optional checksums to detect corruption. Availability is provided through mirroring, clustering, and multipathing. Snapshots provide application-consistent recovery points, while ransomware protection is handled through snapshot rollback and replication strategies, as per my understanding.."

Chathura Nuwan
System Engineer at MIT ESP

Read full review

©2026 PeerSpot, All Rights Reserved

"There are many features that I find very useful, and they are best in industry. First, I will explain from low level to high level. There is a structured feature list. I have always loved the Versus volume manager, and for the volume it has file system on Unix site because it makes storage administration life very easy. It makes it very easy to manage, very easy to deploy, and very easy to recover where needed. I had some use cases during recovery projects where together with very fast support, we were able to recover a completely lost file system and data from scratch without restoring any backup, restoring directly from the volume manager site and from the file system site.

The dynamic multitasking feature is especially useful to have standardized multipetting structures among different operating systems and different solutions, different hardware and software sites. I have loved using latest cluster servers, which you can use for databases, application servers, SQL servers, Oracle RAC, or any application server or anything else you can think of. You can do this on-premise in your own data center or in the cloud without any problems. The volumetric replicator was also something good, though there are some downsides when using replication on some databases, but it is also good.

Although it is not very commonly used right now, when you use InfoScale together with your virtualization structure, it makes it very easy to manage your virtual machines among different data centers or different clusters. You can easily migrate your systems, virtual machines, or clusters from one to another. When you have different storage solutions and hardware storage solutions, and you want to migrate from one storage system to another, such as from EMC to Hitachi or from IBM to something else, you can easily migrate from one storage system to another directly on the fly and online with a few clicks or commands, and it does the job without any interruption on the structures. It is completely online.."

SavasIrez
Professional Services Director – New Technologies at Gantek

Read full review

©2026 PeerSpot, All Rights Reserved

"What I value most about InfoScale is its ease of use and clear visibility into environment operations, particularly in large environments. I set up Veritas Operation Manager, which interacts with all clusters and provides a central management location for the entire clustering environment. The Cluster Manager graphical interface tool is excellent for identifying problems easily. The solution is robust and rarely causes issues. When cluster problems arise, I can identify that the actual problem exists elsewhere, with Veritas alerting me that external factors are affecting cluster behavior. This makes identifying root causes straightforward in a solid environment.

Comparing InfoScale to HACMP (now called PowerHA by IBM), InfoScale is significantly easier. Before 2010 and 2011, my customer used PowerHA before transitioning to InfoScale. We started with version 5 and progressed to version 7.3. The last running instances are version 7.2 or 7.3. We had strong confidence in the product, and architects were very satisfied with its performance matching our requirements exactly. Pricing is the only issue, as the customer cannot justify additional investment and is phasing out instances in favor of custom solutions.

I have used the application-aware failover feature. The environment runs very few single-node systems with application HA, which continue to function but are not actively used. The customer is working to remove applications from Veritas control. This feature remains intact and operational but has no upgrade or evolution plans.

Regarding layering dependencies across web, application, and data tiers, I worked with the Virtual Business Service feature within Veritas Operation Manager. I created Virtual Business Services with dependencies for frontend web servers, databases such as DB2 or Oracle, and defined startup sequences where the database starts first, followed by frontend services only after the database is live and running. This layering system was quite helpful, requiring only a single click to trigger the entire process, leaving Veritas to manage everything automatically. ."

TABISH JAVED
Site Reliability Engineer (Certified) at Kyndryl India

Read full review

©2026 PeerSpot, All Rights Reserved

The best features InfoScale offers include intelligent failover and fallback, along with advanced storage virtualization, which really stand out to me.

Intelligent failover helps my team in day-to-day operations by combining multiple disks into a logical volume and abstracting physical storage from applications. Storage virtualization makes a significant difference for us.

“InfoScale has positively impacted our organization by making our services available to customers at all times. We also have detailed disaster recovery in place, ensuring we are compliant as well.

“I have noticed a reduction in downtime and an improvement in compliance. We are recently compliant to SOC 2 Type 2, where our company is compliant based on this. Using this, we automate processes instead of doing manual documentation, which gives us a complete picture of disaster recovery and response time.

“InfoScale's ability to maintain data integrity and availability during a cyber event such as a ransomware attack ensures that if one system is encrypted, we can fail over to another clean system automatically, preserving data integrity and consistency against malicious activity.

“My experience with InfoScale's application-aware failover feature is the best. Over the last two years, we did not experience any application failover or receive alerts due to the immediate switchover mechanism in active-active mode that ensures no downtime, helping us significantly with confidence and trust in our organization.

“InfoScale's autonomous operational resilience helps significantly reduce system downtime. There has been no downtime for the last two years due to failover, with all disaster recovery processes documented and functioning as expected without any interruptions when switching back to the primary system. .”

Jaswanth Kotla

Sr Network and Security Engineer at a computer software company with 201-500 employees

Read full review

©2026 PeerSpot, All Rights Reserved

“I work primarily with Cluster Server. As an instructor and contractor, I teach the product to other administrators. The biggest part of the product that I have been discussing lately is Cluster Server and the replication, including using the new WORM storage for ransomware protection.

I could discuss the features and capabilities of InfoScale that I have found most valuable for an extensive amount of time, but the general data protection across so many different areas stands out. One does not have to be tied to any particular vendor, and you can make replicated copies over distances with the global option. The ability to use WORM storage for ransomware protection is what I think is the hottest topic in the whole market right now. The ability to work with all different kinds of storage vendors and capabilities is a significant feature at this moment.

I find InfoScale's automated stack-aware recovery feature to be beneficial when recovering from a ransomware event because anything automated is always helpful in a crisis situation. The automated tool is excellent; this is a product that has been proven as I have been working with it for twenty-three to twenty-four years. The company knows what they are doing when it comes to automating and knowing what commands to run for recovery. The simplicity of it in a panic is a significant part of its appeal.

InfoScale's ability to maintain data integrity and availability during a cyber event such as a ransomware attack is excellent. Having the ease of control over the different copies available and being able to bring up another copy within minutes in a safe location on secure storage is invaluable for business continuity. Getting back to business in minutes instead of hours and days is incomparable.

My experience with InfoScale's feature for application-aware failover has been excellent. The heartbeat mechanisms involved in recognizing when a failover needs to take place have even improved with the IMF feature. The system is very quick in recognizing issues; the kernel is the first to recognize a problem with IMF and alerts the agent so action can be taken.

This has been very beneficial for my operations; the availability of everything in

©2026 PeerSpot, All Rights Reserved

quick order is a significant advantage. The monitoring capability allows me to see what is happening with the applications and products overall at any time, which has been excellent.

InfoScale has played a role in reducing downtime, including planned downtime when I can quickly switch over for maintenance. However, unplanned events are where the bigger downtime comes from. The ability for InfoScale to monitor what is happening and react in a timely manner is significant, allowing me to get back to business in minutes rather than hours and days..”

Brad Waddell
Technical Trainer at netxinc

Read full review

©2026 PeerSpot, All Rights Reserved

Other Solutions Considered

"I was previously using a traditional setup for backup and high availability, but I switched to InfoScale because I needed better high availability and automatic failover, improved scalability, and support for critical applications.."

Islam Hamada
Network Security Engineer at EMAK For Integrated solutions

Read full review

"I have experience with Solaris Cluster for high availability, and it is still used in some customer environments based on their requirements. However, for Kubernetes-based environments, we did not previously use a dedicated DR solution. With InfoScale, we are able to address this gap by providing real use cases for ensuring both data and application availability, including automated failover and continuous protection without requiring manual intervention.."

Chathura Nuwan
System Engineer at MIT ESP

Read full review

©2026 PeerSpot, All Rights Reserved

"Comparing InfoScale to other clustering products I have worked with, the heartbeat implementation is distinctive. HACMP does not run a heartbeat; their current setup runs heartbeat on disk. In VCS, I maintained three different types of heartbeats, providing redundancy if one failed. Another excellent Veritas feature is the ability to freeze applications or service groups. Whenever maintenance approached or for other reasons, I could freeze them. This option does not exist in any other solution, making it exceptional. The freeze option is outstanding in Veritas.."

TABISH JAVED
Site Reliability Engineer (Certified) at Kyndryl India

Read full review

"The Cluster Manager graphical interface was excellent for identifying problems easily. The solution was robust and rarely caused issues. When cluster problems arose, I identified that the actual problem existed elsewhere, with Veritas alerting me that external factors affected cluster behavior, making root cause identification straightforward in a solid environment. HACMP, now called PowerHA and developed by IBM, is significantly harder. In my early environment before 2010 and 2011, my customer used PowerHA before transitioning to InfoScale. We started with version 5 and progressed to version 7.3.

Comparing InfoScale to other clustering products, the heartbeat implementation stands out. HACMP does not run a heartbeat; they now have a setup running heartbeat on disk. In VCS, I maintained three different types of heartbeats. If one failed, another remained active. When the second failed, the third provided redundancy. Another excellent Veritas feature is the ability to freeze applications or service groups. Whenever maintenance was scheduled or for other reasons, I could freeze them. This option does not exist in any other solution, making it outstanding. The freeze option is exceptional in Veritas. ."

TABISH JAVED
Site Reliability Engineer (Certified) at Kyndryl India

Read full review

©2026 PeerSpot, All Rights Reserved

ROI

Real user quotes about their ROI:

We have seen a return on investment. We can consider the reduction in reliance on the L1 support team, who generally check analytics of server metrics manually, as a benefit. InfoScale provides automatic remediation steps, saving money and yielding a positive return.

Jaswanth Kotla
Sr Network and Security Engineer at a computer software company with 201-500 employees

Read full review

©2026 PeerSpot, All Rights Reserved

"From my perspective, I see a lot of ROI. As I said, the use of different storage solutions with InfoScale makes your system independent of what storage you are using. Your systems can only see what has been given by InfoScale storage subsystem. You can easily switch from one storage system to another, which is just online and without any impact on the system or most of the time at least. You can actually throttle it.

When you talk about disaster recovery, this is actually something that InfoScale really shines because to have something similar to that, you need to either have different solutions. For virtualization, you will need to perhaps use Zerto or something similar. You will need to use Oracle data guard for Oracle, and you will need to use other solutions for Exchange or SQL or other applications. When you change the storage, you will need to use different storage application solutions. InfoScale does that all by itself without any dependency on different solutions. InfoScale is really product agnostic in that sense. When you have InfoScale, you can actually put almost anything new into your system and use it within your system.

I had an example when I talked about InfoScale with my customers. My example was something to this effect: today, you have a storage solution on fiber channel or maybe network site. But tomorrow, you may have DNA based storage solution, most probably you will be able to use that DNA.."

SavasIrez
Professional Services Director – New Technologies at Gantek

Read full review

©2026 PeerSpot, All Rights Reserved

Use Case

“The use case depends on the client's needs. They have three different components in InfoScale. They have availability, Cluster Server, and the Cluster File System..”

Verified user
Presales Engineer at a tech services company with 51-200 employees

Read full review

“I have worked on implementing various InfoScale products based on customer requirements, including Storage Foundation, High Availability, and Cluster File System to support high availability and workload distribution. I also have experience implementing InfoScale for Kubernetes to support disaster recovery (DR) solutions in Kubernetes environments..”

Chathura Nuwan
System Engineer at MIT ESP

Read full review

©2026 PeerSpot, All Rights Reserved

"To have high availability of data center resources, especially databases and applications, I needed data replicated from one data center to a disaster recovery data center or another data center, and to have disaster recovery using InfoScale features such as global clustering, replication, and other capabilities.

At the beginning, I was a partner, reseller partner and consulting partner at the same time. From time to time, I continued to give customers consultancy around InfoScale where they needed it, but mostly as a reseller and consultant partner.."

SavasIrez
Professional Services Director – New Technologies at Gantek

Read full review

©2026 PeerSpot, All Rights Reserved

We primarily use InfoScale for redundancy, specifically high availability for critical patient servers running on the cloud.

A specific example of how we use InfoScale with our patient servers is that it automatically detects failures and switches over the network, applying the nearest load balancer, so only the secondary system gets up.

“Failover for our main use case occurs automatically, and failover takes place to the next server.

“We have deployed InfoScale on-premises.

“We have used AWS in conjunction with InfoScale. .”

Jaswanth Kotla

Sr Network and Security Engineer at a computer software company with 201-500 employees

Read full review

©2026 PeerSpot, All Rights Reserved

“I use InfoScale, and the use case depends on the client's needs. InfoScale has three different components: availability, cluster server, and cluster file system.

I used InfoScale for high availability and automatic failover to minimize downtime. If one server goes down, another takes over automatically. The application stays running without interruptions and data remains consistent and accessible. I also used it in cluster server and cluster file system.

InfoScale is not being used for ransomware protection. I use it only for failover, cluster server, and cluster file system at the application layer.

This provides high availability of data center resources for the database and application. I needed data replicated from one data center to the disaster recovery data center, and I have disaster recovery using InfoScale.

I use high availability to reduce downtime. If one server goes down, another takes over automatically.

If your environment is critical, you should use InfoScale. However, if it is small or non-critical, it is not a good choice because the licensing cost is very high.

InfoScale is good for high availability and high availability clustering. I recommend it for this purpose..”

Islam Hamada
Network Security Engineer at EMAK For Integrated solutions

Read full review

©2026 PeerSpot, All Rights Reserved

“ My work focused on a high-availability environment where customers maintained two or three data centers designed for disaster recovery solutions. I managed local clusters as well as global clusters, and when a service failed in a particular environment, it automatically moved to a different region. The entire solution was designed with InfoScale at its core.

I find InfoScale's automated Stack Aware Recovery feature to be very beneficial when recovering from ransomware events. The system automatically moved services to alternate nodes, behavior I observed while running on UNIX systems that had underlying issues. Veritas was very active and proactive, notifying me in advance about cluster conditions and recommended actions. From a recovery perspective, my main goal was to establish application recovery, which is why I selected this solution. Some critical machines in the environment still run Veritas because it excels at recovery, is easy to understand, and facilitates root cause analysis. Most of the root cause analyses I performed over ten years were not related to Veritas. I consider it a great product that meets expectations. The primary concern is licensing cost, as the customer is unwilling to invest further and has begun cost-cutting measures. With cloud adoption, they are moving workloads to the cloud, believing it offers greater benefits than on-premises solutions. All VCS instances and Veritas clusters ran on-premises only, with nothing moved to the cloud. Most licenses have already expired, and the customer has allowed me to continue using them while exploring alternative solutions. The application team is redesigning applications from scratch, with several already migrated to microservice architecture in Kubernetes in the cloud.

From a recovery standpoint, InfoScale is excellent and easy to manage. A single technical person can handle 100 machines or one application spanning multiple clusters. I utilized the Virtual Business Services feature to design the solution, enabling all databases, frontend machines, backend machines, and related components to move to different clusters seamlessly without any issues.

InfoScale has significantly reduced downtime for my customer. I encountered unusual split-brain issues. Because I did not utilize all cluster features such as I/O Fencing, which requires additional setup and licensing cost, the solution was not

©2026 PeerSpot, All Rights Reserved

designed with I/O Fencing. When split-brain occurred, I had to investigate the cause, protect the data, and determine remediation steps. For data protection, I implemented SCSI-3 Persistent Reserve at the storage level instead of using I/O Fencing. ."

TABISH JAVED
Site Reliability Engineer (Certified) at Kyndryl India

Read full review

©2026 PeerSpot, All Rights Reserved

Setup

The setup process involves configuring and preparing the product or service for use, which may include tasks such as installation, account creation, initial configuration, and troubleshooting any issues that may arise. Below you can find real user quotes about the setup process.

My experience with pricing, setup cost, and licensing indicates that pricing is a little higher and should be reduced since most companies cannot afford it.

Jaswanth Kotla

Read full review

Sr Network and Security Engineer at a computer software company with 201-500 employees

"I called the vendor to set it up, and I observed the installation and the basic configuration. It is easy to integrate with the server and has clear steps for clustering and service setup. The setup is very easy.."

Islam Hamada

Read full review

Network Security Engineer at EMAK For Integrated solutions

©2026 PeerSpot, All Rights Reserved

"The initial setup and deployment of InfoScale is quite smooth. The installer has been good all along and has even improved over time. It is very efficient at understanding what you need, and as long as you are aware ahead of time of requirements such as cabling, system names, and the communication, the installation and configuration process is straightforward and quite smooth.."

Brad Waddell

Technical Trainer at netxinc

Read full review

"My experience with the initial setup and deployment of InfoScale involved implementing it in a traditional method on-premises using Storage Foundation and High Availability and Cluster File System, requiring physical layer setup, VLAN assignments, and package installations across servers with specific configurations for disk-based and majority-based fencing.

In terms of initial setup and deployment ease, the documentation provided is clear, making it easier to set up, although troubleshooting can be difficult due to extensive logs that require deep dives.."

Chathura Nuwan

System Engineer at MIT ESP

Read full review

©2026 PeerSpot, All Rights Reserved

"I was involved in the deployment process and initial setup of InfoScale. We designed the solution, not only myself but also an architect. My role as an SRE involves running the service and maintaining uptime rather than participating in design and solutioning. In the initial days, I set up cluster servers from scratch, installing all LPARs on AIX and configuring storage, then loading Veritas. One important Veritas feature I should mention is that it has maintained hybrid solutions since 2006 or 2007. Veritas can run two different cluster nodes with two different operating systems, such as one Windows node and one Linux node. This feature does not exist in any other solution and is impossible elsewhere. Veritas offers this capability, which I appreciate. I never utilized this feature, but its availability demonstrates Veritas' comprehensive approach.."

TABISH JAVED

Site Reliability Engineer (Certified) at Kyndryl India

Read full review

©2026 PeerSpot, All Rights Reserved

"The actual setup involves planning installation on a new system, then planning and configuring the services on that new system and understanding how to move services from production to a new InfoScale system. Then copying production data to the new system and testing it together with customers. At some point, you bring down the production and bring up the new system and test it on that new system to ensure everything is working fine, and then configuring the dependencies with other systems if needed and connecting all these services from one to another.

This actually depends on where you are implementing InfoScale. If you are implementing on a database system, you need to have a good and structured checklist to understand if you need to migrate or if you need to do it on the same system, a conversion. In general, when I scale InfoScale on a data center for a customer, I usually prefer to do it implementing on new systems and moving data and services from one system to another and checking that everything is alright and can be brought up.."

SavasIrez

Professional Services Director – New Technologies at Gantek

Read full review

©2026 PeerSpot, All Rights Reserved

Customer Service and Support

“I often communicate with the technical support of InfoScale.

In my impression, InfoScale's support specialists are knowledgeable, providing remote and ticket-based assistance while giving quick solutions and thorough documentation; I believe they offer better help..”

Chathura Nuwan
System Engineer at MIT ESP

Read full review

“When I encounter situations I cannot resolve or understand, or when incidents require vendor input or investigation, I contact Veritas customer service. I raise tickets, and they participate in root cause analysis and incident fixes. This interaction is limited because the product is stable and robust, rarely causing problems. Once the Veritas InfoScale contract expired, my customer designed alternative solutions outside InfoScale. They began phasing out and decommissioning InfoScale environments, reducing from 500 cluster nodes to approximately 50 nodes. Product interaction with the support team is now very limited..”

TABISH JAVED
Site Reliability Engineer (Certified) at Kyndryl India

Read full review

©2026 PeerSpot, All Rights Reserved

“I have communicated with InfoScale's technical support and customer service.

My experience with the technical support specialists has generally been very good. When I am on-site, I sometimes get anxious because I want to get the customer back in business quickly. If I do not see the issue right away, I open a support ticket, but nine out of ten times, I find the problem myself before the support team gets back to me. They have always been very helpful when I have needed them. However, I often open tickets prematurely just because I want to expedite the process for the customer.

Based on my experience with support, I would rate them a nine, only because occasionally the first person I talk to does not know more than I do and it needs to be escalated to reach someone more knowledgeable. These cases typically involve more difficult problems, so I still rate them very high for their support..”

Brad Waddell
Technical Trainer at netxinc

Read full review

©2026 PeerSpot, All Rights Reserved

"It is always important to think about the dependencies on your systems. Consider this picture: you have an application that needs to have a back UI and an application site and different APIs. They also need to connect to one or multiple databases underneath. When you want to have high availability and easy recoverability from a disaster recovery perspective or from a high availability perspective, it is very important to have the system fail over from one to another with a strict definition of dependency.

You need to have a definition of dependency, and that dependency needs to have a correct order. For example, if you do not have the database open, you cannot bring up your application. From high availability or disaster recovery perspective, if you have planned failovers, you will need to bring down services one by one from the upper side going to lower according to your dependency chart. You need to bring down services on the UI, then application. Right after everything is done, you can bring down your database and switch over to a different site. Then you need to have the database come online, and only after the database is online, depending on your need for dependency, you should bring up your application. InfoScale does this job very well in a manner that it can automatically bring down services one by one on one side and bring them up on another. When you see every system, every application, and service online on a site, you can be sure that everything has been brought up according to your dependency charts.."

SavasIrez
Professional Services Director – New Technologies at Gantek

Read full review

©2026 PeerSpot, All Rights Reserved

Other Advice

My advice for others looking into using InfoScale is to focus on automated redundancy and high availability. I would rate this review as an eight out of ten.

Jaswanth Kotla
Sr Network and Security Engineer at a computer software company with 201-500 employees

Read full review

"This product works very well, but prior to implementing, you need to study how the application works. Like all software, you must be follow all the recommendations from the manufacture. ."

Verified user
Storage and VMware Expert at a comms service provider with 1,001-5,000 employees

Read full review

©2026 PeerSpot, All Rights Reserved

"This part is sort of covered during these last two questions because think of it this way. When I worked on a disaster recovery project, we had to test the disaster recovery. All I did during the disaster recovery test was to switch over from one first primary site and every service on the primary site and switch everything to the other side. This has been done during all these features together, and everything is brought down on one side and brought up on the other side. This is actually done by all these features altogether using all these features. My overall rating for this product is nine point five.."

SavasIrez
Professional Services Director – New Technologies at Gantek

Read full review

"I recommend InfoScale to others based on my extensive experience. Previously, when advising another customer interested in IBM solutions such as PowerHA and HACMP who had purchased an IBM solution, I suggested they select Veritas instead. Despite being an IBM employee at that time, I recommended they not purchase PowerHA and proceed with Veritas. They discussed my explanation and decided to trust my assessment, reasoning that managing two different systems would be complicated and that my comfort with Veritas made it the better choice. ."

TABISH JAVED
Site Reliability Engineer (Certified) at Kyndryl India

Read full review

"I have mainly used InfoScale with Oracle Database on Solaris systems and application-aware failover. My advice would be to properly configure the Oracle agents and resource groups, as this ensures fast failover and reduced downtime,

©2026 PeerSpot, All Rights Reserved

typically around one minute.

“Based on everything I have shared on various aspects of InfoScale, I would rate it 8.5 out of 10..”

Chathura Nuwan
System Engineer at MIT ESP

Read full review

©2026 PeerSpot, All Rights Reserved

“I have not been in a situation where we have had serious problems; we have always gotten ahead of issues before they arise. The biggest tangible benefit comes from having the knowledge of how long it takes to rebuild from a disaster. People reach out to us because they want to avoid those problems, and I have successfully helped them with this.

The biggest challenges I have faced over the years mostly come from the training side, explaining some of the under-the-covers features so that people understand how it works. I remember twenty-some years ago hearing it for the first time and wondering if I would ever understand what it meant. The longer you work with something, the more natural it becomes, allowing me to explain it to new users better. The terminology can vary between different products, which sometimes complicates understanding.

My customers usually prefer a combination of deployment options for InfoScale; more and more are moving to the cloud rather than maintaining a physical DR location. They typically start on-premises but with the objective of replicating or duplicating to the cloud for security reasons.

When it comes to cloud preferences, I refer to the big three regarding the formerly known Veritas products; Amazon, Google, and Azure cloud are the most common choices among my customers.

I have been impressed with InfoScale from the very beginning all those years ago. The ability to handle array failures, whether it is just a disk or connectivity issues to an array, has been very beneficial. The failover features and clustering capabilities have been incredible; another big part of it is not being tied to a vendor. If we have an office in a different city or country that is getting a better deal on hardware, it does not matter what they are using. InfoScale communicates with all of them and allows for data migration and replication, providing significant flexibility.

I have been researching it and looking into it, and it looks excellent. I have also worked with all of the Veritas products in the past, so when it came to recovery, I

©2026 PeerSpot, All Rights Reserved

was more focused on the NetBackup area. The fact that we can have a copy that is protected and replicated quickly is more beneficial than recovering from a backup.

I give this review a rating of ten out of ten..”

Brad Waddell
Technical Trainer at netxinc

Read full review

©2026 PeerSpot, All Rights Reserved

Top Industries

by visitors reading reviews

Financial Services Firm	20%
Manufacturing Company	8%
Computer Software Company	8%
University	7%

Company Size

by reviewers

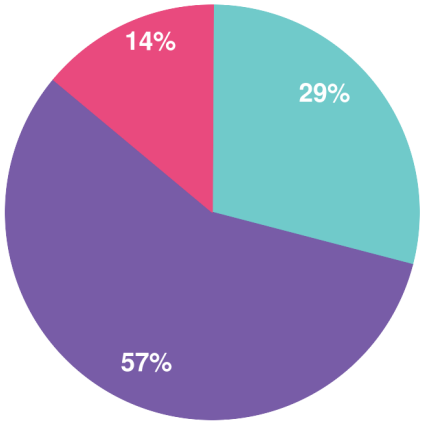

by visitors reading reviews

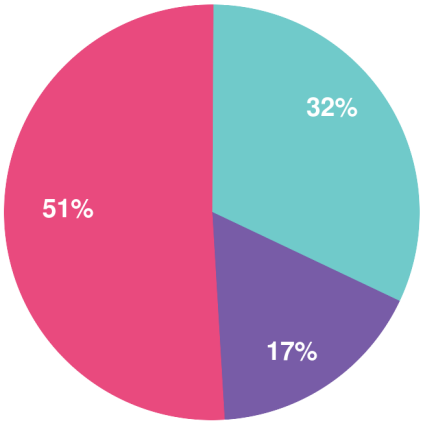

Large Enterprise Midsize Enterprise Small Business

©2026 PeerSpot, All Rights Reserved

About this buyer's guide

Thanks for downloading this PeerSpot report.

The summaries, overviews and recaps in this report are all based on real user feedback and reviews collected by PeerSpot's team. Every reviewer on PeerSpot has been authenticated with our triple authentication process. This is done to ensure that every review provided is an unbiased review from a real user.

Get a custom version of this report... Personalized for you!

Please note that this is a generic report based on reviews and opinions from the collective PeerSpot community. We offer a customized report of solutions recommended for you based on:

- Your industry
- Company size
- Which solutions you're already considering

The customized report will include recommendations for you based on what other people like you are using and researching.

Answer a few questions in our short wizard to get your customized report.
Get your personalized report here

©2026 PeerSpot, All Rights Reserved

About PeerSpot

PeerSpot is the leading review site for software running on AWS and other platforms. We created PeerSpot to provide a trusted platform to share information about software, applications, and services. Since 2012, over 22 million people have used PeerSpot to choose the right software for their business.

PeerSpot helps tech professionals by providing:

- A list of products recommended by real users
- In-depth reviews, including pros and cons
- Specific information to help you choose the best vendor for your needs

Use PeerSpot to:

- Read and post reviews of products
- Access over 30,000 buyer's guides and comparison reports
- Request or share information about functionality, quality, and pricing

Join PeerSpot to connect with peers to help you:

- Get immediate answers to questions
- Validate vendor claims
- Exchange tips for getting the best deals with vendor

Visit PeerSpot: www.peerspot.com

PeerSpot

244 5th Avenue, Suite R-230 • New York, NY 10001

reports@peerspot.com

+1 646.328.1944

©2026 PeerSpot, All Rights Reserved