
FAQ's

Why would I use GeoCluster instead of setting up Microsoft Clusters (MSCS) with shared disk resources?

In an MSCS cluster the shared disk subsystem and storage path (SCSI or Fibre Channel) represents a single point of failure for the entire cluster. If the power goes out on the disk subsystem or someone knocks the SCSI cable out, the cluster will go down. Both the cluster nodes and the shared storage are vulnerable to geographic disasters such as fires, flood, earthquakes or hurricanes. Since GeoCluster utilizes independent storage devices on each member node and replicates data between nodes over existing network links using TCP/IP, you are not susceptible to downtime due to a disk failure. And since GeoCluster removes the distance limitations between nodes, you can protect yourself against regional and geographic failures by separating the cluster nodes into different geographic locations.

MSCS, in a shared disk configuration, requires you to use Microsoft certified hardware configurations of servers, controllers, and disk subsystems. Typically, you cannot take a standalone server you have running and decide to cluster it using MSCS alone. GeoCluster can be installed on any 2000/NT supported server and only requires that disks exist on each server with matching drive letters and the same amount of disk space on the partitions that will be replicated, for both nodes of the cluster.

If I want to provide high-availability or "failover" for an application, should I use GeoCluster or Double-Take?

In most cases you can use either product. The best answer will depend on several factors and your objectives.

Server Failover vs. Application Failover - Double-Take provides the ability for a target server to monitor one or more source servers and manually or automatically assume the identity of the source server, if it fails. Failover includes assuming the name, IP address, and shares and the ability to run scripts to automate other tasks such as restarting services or sending notifications. In order to bring the original source machine back online it is necessary for the failback process to release the source identity from the target. You may also need to restore the latest data from the target server back to the production server to resume operation.

If you have multiple applications running on the same server or want a tighter level of integration with the applications, GeoCluster should be considered. GeoCluster can provide application monitoring and failover of single or multiple applications, instead of the entire server.

Virtual Servers - In a GeoCluster, the physical identities of the nodes do not change but each application can create a virtual server that end users see. Each virtual server can be easily moved back and forth between nodes in the cluster without impacting the users view. A single cluster may contain many virtual servers that can be moved independently. This approach eliminates the need to failback and restore data to the original server after a failover. Virtual servers in GeoCluster also make active/active clusters simple to configure where each geographically dispersed node backs up the other. For example, Microsoft SQL Server™ might be running on Node 1 while Exchange is running on Node 2.

With Double-Take the actual server identity is assumed by the target server. The concept of having virtual servers is not inherent to Double-Take.

Cluster Awareness - If an application is available in a MSCS cluster aware version

(sometimes called the Enterprise Edition), that version of the application is designed and tested by the original developer to be monitored by and moved between cluster nodes as needed. Configuration changes that are stored in the registry are made available to all nodes eliminating the need to manually match configuration changes. Some cluster aware applications provide special capabilities such as the ability to run multiple instances of an application on different nodes within the cluster and move those instances individually, or the ability to automatically install required components to all nodes within the cluster. To take advantage of applications that were written to capitalize on Microsoft Cluster Server, you should use GeoCluster for maximum performance and protection of your applications. Either product can generally be used to provide failover of noncluster aware applications or services through generic mechanisms.

Hardware Flexibility - While GeoCluster clusters require matching drive letters for replication, standard Double-Take allows virtually any number of machines to be replicated to a single target server and for a single target to assume multiple names and IP address simultaneously. Drive letters and sizes do not need to match and Double-Take automatically adjusts shares to point to the proper location on the target server. Double-Take replication sets allow directories and/or individual files to be included or excluded from replication allowing greater control whereas GeoCluster replicates full drives.

Replication Options - While the core replication technology is the same, Double-Take offers additional options and flexibility not provided by GeoCluster. These include:

File Level Selection - Include or exclude directories, files, and/or wild card specifications Multi-Volume Replication Sets-Data from multiple volumes can be defined as a single replication set

Target Path Control - Double-Take allows virtually unlimited control over the target drive and path where replicated data is stored (GeoCluster only replicates "drive to drive", i.e. d: to d:, root to root

Any-to-Any Connections - Double-Take can replicate in virtually any configuration including one-to-many, many-to-one or many-to-many with no limitations on network location, domain or OS type. You can for example replicate from a workstation to a server or between servers in trusted or non-trusted domains or machines on different IP subnets. GeoCluster is subject to configuration limitations imposed by MSCS such that cluster nodes must be members of the same Windows® Domain and have network connections on the same IP subnet (although this can be accomplished remotely using VLAN technologies).

Does GeoCluster support clustering multiple instances of Microsoft SQL Server™ on a cluster?

GeoCluster supports the same application configurations as MSCS allows on a shared disk configuration. Multiple SQL virtual servers can be created in multiple cluster groups and moved independently between each node.

If I want to provide data protection or off-site disaster recovery for an application, should I use GeoCluster or standard Double-Take?

Both products provide real-time data protection by maintaining multiple replicas of critical data on multiple machines. If your goal is simply to provide data replication for disaster recovery purposes of either a standard server or an existing MSCS cluster, you should use Double-Take. The main purpose of GeoCluster is to remove the MSCS limitation of a shared disk and allow the nodes of a cluster to be geographically dispersed. GeoCluster does not currently provide all the advanced options of Double-Take such as individual file and directory selection

or wild card include/exclude rules. These options provide optimized file level data replication. Optimal data transfer is beneficial in limited bandwidth environments. The decision really comes down to your applications, business needs, and goals. In cases where you want high availability AND disaster recovery, you can use both GeoCluster and Double-Take together, which is offered by NSI as GeoCluster Plus, combining the benefits of both solutions into a single product offering. For example; you could split the nodes of your cluster and place them in different builds on the same campus. This would provide high availability for your applications using the GeoCluster portion. You could also replicate the application data to a different part of the country using the Double-Take portion. GeoCluster Plus combines the best of both worlds.

Double-Take runs on NT, Windows 2000, Windows Server 2003, and Windows Storage Server 2003. What about GeoCluster?

GeoCluster is an extension of Microsoft Cluster Server and can be used with operating systems that support Microsoft Cluster Server. This includes Windows NT 4.0 Enterprise Edition, Windows 2000 Advanced Server or Datacenter, and Windows Server 2003 Enterprise and Datacenter. GeoCluster cannot be used on standard Windows NT Server (non-Enterprise Edition), Windows NT Terminal Server Edition, Windows NT Workstation, Windows 2000 Professional, Windows 2000 Server, or Windows Server 2003 Standard or Web Editions. Most of these operating systems provide easy and affordable upgrade paths to an operating system that includes MSCS and in most cases you do not need to update your CALs (Client Access Licenses). Check with Microsoft on upgrade options and pricing.

Does GeoCluster provide any load balancing functions?

GeoCluster takes advantage of the load balancing capabilities provided by MSCS which are generally limited to balancing the distribution of cluster groups among the nodes. In other words, instead of running SQL, IIS, and Exchange on a single node while the other nodes remains idle, it may be preferred to move Exchange to the second node to even the processing load. GeoCluster can also be used on a MSCS cluster in conjunction with Microsoft Windows Load Balancing Service (WLBS) which provides IP load balancing to multiple servers within a cluster (which would contain read-only replicas of the data).

How does GeoCluster or Double-Take ensure data integrity between each replica?

The replication engine used by both Double-Take and GeoCluster captures byte-level changes being made to monitored, volumes, directories, or files and transmits those changes to the target server(s) in exactly the same sequence they originally occurred to maintain consistency of the data at all times. If the source data is changing faster than it can be transmitted, the data is queued and immediately transmitted in time sequenced order once bandwidth is available.

Does that mean the replicated data would be old?

Data is old only if it is being changed faster than it can be transmitted to the other node. Under normal conditions, data would be transmitted and applied immediately. Even if data is queued up for minutes or hours due to a network outage, for example, the data on the target represents a consistent image of exactly what the source looked like at the time the last replication packet was received. If the replicated data needed to be used at that point, the applications would respond exactly as if the source machine had been powered off at that point in time and restarted. Most applications will perform their own consistency check and roll back any incomplete transactions and resume normal operations from that point. This time difference is referred to as "latency" or "in-flight data" and should not be confused with data

integrity or consistency. GeoCluster provides SNMP/WMI monitoring functionality which can be used to monitor the latency of the target data set which allows pro-active network outage resolution.

How can I verify that the data matches? What happens if they aren't synchronized, for example, when a node is down for a long period of time?

GeoCluster provides a verification option that can be used to compare the replicas at any time or on a scheduled basis. This process compares each file and can even compare individual blocks of large files for differences, and it can automatically correct any discrepancies. When a node has been down and then rejoins the cluster, GeoCluster will automatically begin a block checksum comparison of the data on all possible owning nodes. It will automatically resynchronize each node. Until this resynchronization is complete, the group(s) are prevented from being moved to any node that may contain inconsistent data.

If I create a replicated disk resource for drive D: and it is currently online on Node 1, what will I see if I look at the D: drive on node 2? Can I use this copy of the data? Could I accidentally change or corrupt the replica (copy of the data on the non-owning nodes)?

Drive D: is mounted and would be visible locally on node 2 but would be a read-only device to prevent any accidental corruption. You could, however, use this replica for read-only purposes such as performing tape backup or even for a web cluster using Windows Load Balancing Service.

How far apart can the nodes be? Are there any limitations on the network configurations allowed?

There is no inherent distance limitation between the nodes in a GeoCluster, however there are some restrictions within MSCS that could limit distances. Currently the network connections between each node must be on the same IP subnet. For example, each node may have both a private network connection (192.168.1.1 and 192.168.1.2 on 255.255.255.0) and a public connection for user access (10.0.0.1 and 10.0.0.2 on 255.0.0.0). The NICs must be on the same IP subnet because the cluster heartbeat messages are IP broadcasts and will not cross routers. This configuration could be established across great distances using VLAN (virtual LAN) and VPN (Virtual Private Network) capabilities in many modern networking devices. MSCS has imposed a maximum 500ms response time for heartbeats between cluster members. As GeoCluster runs in conjunction with MSCS, the servers must be able to maintain this response time. This is important to consider when considering networks and locations for each cluster member.

Why does GeoCluster provide compression as an option and not just have it always on?

While compression of transmitted data can have a very positive impact on the data protection solution, it is not always optimal to utilize it, for a few different reasons. Compression requires additional processing which results in slightly higher CPU usage. While some servers and applications can allow for this, there may be some where this is not acceptable. For this reason, NSI chose to allow its customers to determine when and where compression is suitable for their business and their data. GeoCluster not only allows users to configure compression for on or off, but also allows users to set compression to one of three different levels, depending on the particular data being replicated. As the leader in enterprise data protection solutions, NSI Software provides the flexibility to meet the varying needs of its customers by allowing them to choose the level that is right for them.

What methods of notification and alerting are available within GeoCluster?

GeoCluster provides numerous methods of notification for events occurring in the replication and data protection environment. Events and messages can be viewed via the Windows Event Viewer as standard Windows events; can be viewed via a framework or other SNMP-compatible source via our SNMP traps, or via standard email, which could include pagers and cell phones enabled for SMTP. Separate notification methods and recipients can be configured on each node of the GeoCluster.

Can GeoCluster restrict administrators to only their set of servers?

While GeoCluster does not prevent administrators from viewing or managing certain servers, it does allow individual administrators to create custom views of just a particular set of servers. Instead of searching through tens or hundreds of servers in the management console looking for a particular one or few, administrators can create a view that will display just the server(s) that they are concerned with. Scalability is not an issue with GeoCluster, as hundreds of servers can be viewed in smaller subsets for separation of tasks between administrators and for ease of management.

What is Majority Node Set from Microsoft?

Majority Node Set is the Microsoft solution to the same problem that our arbitration resolved and is available in Windows Server 2003. Since Microsoft has expanded the number of possible nodes within a cluster they had to provide a way to prevent split brain (each cluster member thinks it owns the same resource) scenarios. In a regular (old) cluster model, the quorum was merely a file on a dedicated volume that maintained a listing of which node owned which resource. Since we split the physical media, this meant the quorum was now located in 2 different locations and could potentially allow each node to think it was the owner (split brain). To prevent this, we force arbitration. Arbitration is the process of contacting several points/servers to determine who really owns the resource. If it cannot gain a majority vote from the multiple points, it cannot become owner of the resource. This model allowed us to easily implement more than 2 nodes in a cluster. MS now does something similar with Majority Node Set allowing multiple nodes within a cluster. We recommend customers use Majority Node Set in clusters with more than two nodes.

Does the Majority Node Set feature within MSCS replace the need for GeoCluster?

Majority Node Set does not provide the functionality and benefits that are provided by GeoCluster. Majority Node Set works in concert with GeoCluster. Within MSCS there is a need to maintain quorum devices in order for it to know when a cluster member owns a particular resource. Because GeoCluster splits up the storage resources unbeknownst to MSCS, GeoCluster would have to perform additional processing and replication in order to provide the quorum resources needed by MSCS. With the addition of Majority Node Set in MSCS, the need for GeoCluster to maintain these quorum devices is removed. MSCS now can be configured to maintain quorum on it's own even when using non-shared storage devices. This removes additional work from GeoCluster and allows the clustering application to handle this on it's own. This feature of MSCS allows for a much tighter integration with GeoCluster. GeoCluster is the only product on the market to take advantage of Majority Node Set.

Is GeoCluster Windows certified?

GeoCluster has been certified on Windows 2000 Advanced Server and Datacenter and on Windows Server 2003 Enterprise and Datacenter. Datacenter level certification requires

stringent levels of stability and performance.

Why is GeoCluster not certified on Windows 2000 Server Edition or Windows Server 2003 Standard Edition?

GeoCluster requires Microsoft Clusters (MSCS) as a prerequisite, as GeoCluster is an extension of MSCS, improving its level of protection. As MSCS is not available on Windows 2000 Server Edition or Windows Server 2003 Standard Edition, GeoCluster inherently cannot run on these versions, therefore it is not possible to certify on them.

Should Majority Node Set always be recommended when running GeoCluster?

The Majority Node Set (MNS) quorum is not something recommended in all scenarios. In cases of two-node clusters, the MNS quorum is not reasonable because failure of one node will mean the remaining node cannot arbitrate with a majority (>50%) and will shut down the cluster service. With an MNS quorum a majority of the nodes must be present in order for the cluster to come online $((\text{number of nodes} / 2) + 1)$. If it is already online, the cluster will go offline if that many nodes are not available.

It is probable that customers will not be using the MNS quorum until they have 5+ nodes in the cluster. With four nodes, you can still only lose one node (you have to have three up to have a majority)--with five nodes you can lose one or two nodes without bringing the cluster down.

Does MNS compete with or replace Double-Take or GeoCluster?

MNS does not perform replication nor does it allow a MSCS configuration to utilize non-shared storage. It can simplify the work required by GeoCluster in that it will handle the cluster quorum instead of GeoCluster having to handle it and control resources.

Does Microsoft Cluster Service need to be installed separately from GeoCluster?

Yes, MSCS must be installed on the server prior to installing GeoCluster. The GeoCluster installation process checks for MSCS and if it is not configured the install will not continue.

Is the user required to use both the MSCS interface and the GeoCluster interface to manage their cluster?

GeoCluster is integrated with the MSCS cluster administrator interface allowing the cluster to be managed from within one interface. There is not a specific GeoCluster user interface. Tasks are handled from within the MSCS Cluster Administrator and the NSI Management Console.

Tasks performed from within the MSCS Cluster Administrator:

- Creating the replicated disk resource
- Bringing the replicated disk resource online
- Taking the replicated disk resource offline
- Modifying the replicated disk resource properties

Tasks performed from within the Double-Take Management Console :

- Verifying the data is identical on each node
- Initiating a mirror for establishing a connection
- Moving or removing orphan files
- Changing Double-Take program settings

-
- Monitoring Double-Take activity

Is specific testing, configuration, or integration work necessary for running GeoCluster with particular applications such as Exchange or SQL Server?

No, there are no specific certifications or configuration changes required to run GeoCluster with any layered applications. Because GeoCluster sits on top of MSCS, by default it supports any application and/or configuration certified by MSCS. GeoCluster does not interact directly with the applications and therefore does not need specific testing or engineering to work with them.

What if a cluster resource moves to a different node while replication data is still in the queue?

GeoCluster will pause the process of the resource moving to another cluster node until it has flushed the queue of any data to be replicated. This ensures that the resource on the other cluster nodes is current.

. Privacy Policy