the Availability Digest

www.availabilitydigest.com

Neverfail for Windows Applications

June 2010

Neverfail, from Neverfail Ltd. (www.neverfailgroup.com), ensures continuity of user services provided by Microsoft Windows applications via data replication and automated failover procedures. Neverfail maintains a local or remote standby server in complete application synchronization with its primary companion via asynchronous data replication. It monitors the health of an entire application ecosystem at the business level. If it cannot resolve a potential availability issue detected on the primary server, Neverfail will failover the application to the standby server.

Overview

Neverfail is focused on Windows applications. Application monitoring and failover polices are determined by rules configured into Neverfail. For many common applications, such as Exchange, SQL Server, File Server, IIS, SharePoint, and BlackBerry Enterprise Server (BES), rules are preconfigured into plug-ins available from Neverfail. Users may modify these rules if desired and may establish rules for applications not supported by Neverfail plug-ins.

Neverfail provides a shared-nothing architecture that supports both physical and virtual environments. It is hardware agnostic, though the software must be identical on both servers. Neverfail automatically detects software changes on the primary server and replicates these changes to the standby server. This eliminates “configuration creep” caused by configuration changes in the primary server that do not get reflected in the standby server.

Not only does Neverfail eliminate unplanned downtime, but planned downtime is eliminated as well by upgrading the standby server and then switching over to it while the primary server is being upgraded.

Neverfail’s replication engine and failover procedures can satisfy short RTOs (recovery time objectives) and RPOs (recovery point objectives). Failover to a standby server following a primary-server failure is typically about two minutes (RTO). Data loss following a primary server failure (RPO) is measured in seconds if the standby server is collocated and in minutes if it is remote. User sessions are maintained during a failover so that users do not have to log on again.

Ecosystem Management

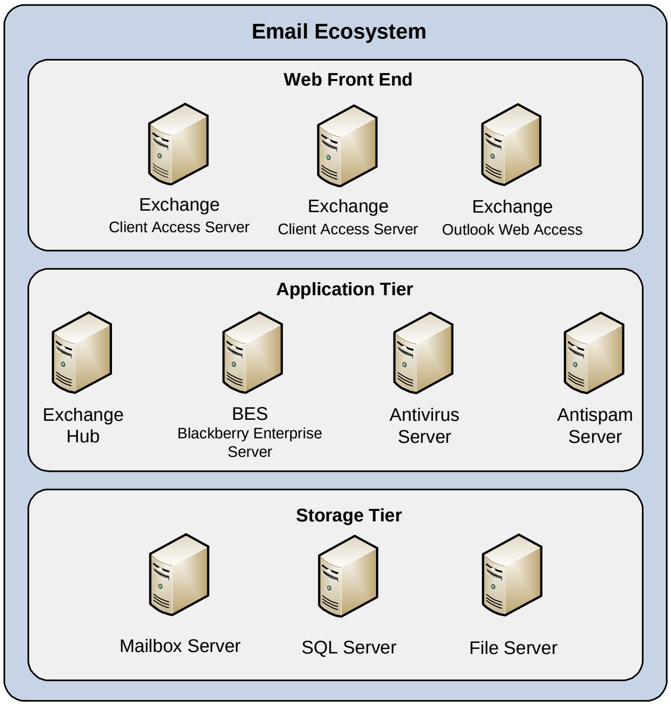

Neverfail does not just protect a process, a database, or an application. It protects the ecosystem for an entire business service.

An ecosystem comprises all of the servers, storage units, applications, data, configuration parameters, networks, and other components that are required to provide a particular service. Ecosystems can be defined based on any criteria, such as technology (email), business function (e-commerce), geographical (New York, London), or any other desired view.

For those functions for which Neverfail provides a plug-in, the plug-in specifies the components of its ecosystem and the rules for monitoring and recovering the ecosystem. For other applications, the user defines the components and rules making up his custom ecosystem via the Neverfail Continuous Availability Director.

According to the failover policy that has been established for an ecosystem, a failure in one of its components might cause just that component to fail over, multiple components to fail over, or the entire ecosystem to fail over. For instance, if failover is to a remote site that is thousands of miles away, it may be better for performance reasons to fail over the entire ecosystem to minimize communication latency between the ecosystem's components.

Architecture

System Configuration

A Neverfail environment includes a primary server and a standby server. The servers may be collocated in the same data center, or they can be hundreds or thousands of miles apart to provide disaster tolerance.

Each server is connected to two networks:

- A private network for heartbeats and data replication.
- A public network for user access.

The private network is extremely critical, as without it the standby server cannot be kept synchronized with the primary server and is useless as a backup. Neverfail supports an optional dual private link to ensure server interconnectivity. This is especially important if the servers are remote from each other and communicate over a WAN.

Neverfail supports heterogeneous hardware configurations. The hardware used by the two servers does not need to be identical.

However, the software on the two servers must be absolutely identical. The operating systems on each machine must be the same and must be at the same service pack and patch level. Disk names, directory structures, and database schemas must be identical. All application executables must be the same version. User authorizations must be identical. Both servers must have the same name, and both must use the same IP address for the public network. The latter requirement allows users to be moved from the primary server to the standby server without losing their connections.

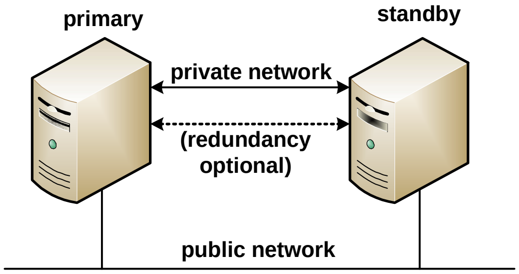

Neverfail Primary/Standby Pair

Neverfail ensures that the initial software installations are identical, and it discovers any changes made thereafter to the primary server. The changes are replicated to the standby server, thus ensuring software homogeneity.

Neverfail supports Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2.

The Application-Aware Management Framework (AMF)

The heart of Neverfail is the Application-Aware Management Framework, AMF. AMF is responsible for Neverfail's core functions:

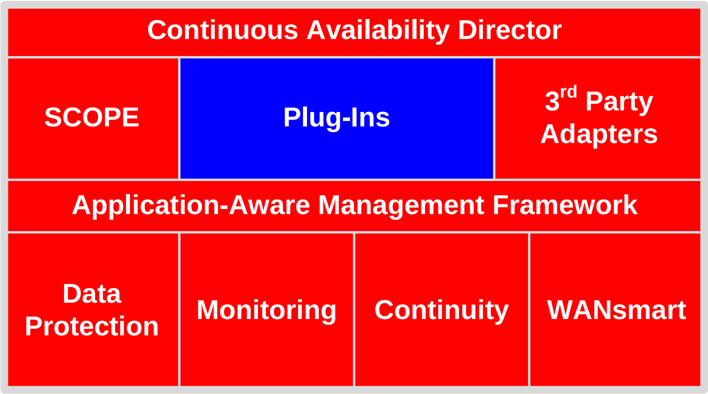

Neverfail Architecture

<u>Data Protection</u>

The Neverfail replication engine maintains real-time consistent copies of all critical data by replicating it asynchronously from the primary server to the standby server. This includes all data that is required to ensure that the standby applications are true clones of the primary applications:

- Databases and file systems
- Database schemas
- Application executables
- Registry settings
- Application settings

Neverfail replicates not only changes to files and tables but also any change to the application configuration. In this way, it is known that the configuration of the standby server is always the same as the primary server. This is necessary to ensure that the standby server will perform properly if it is put into operation as a result of a failover.

Since Neverfail uses asynchronous replication to replicate data and configuration changes, some data may be lost following the failure of the primary server. If the primary and standby servers are collocated, this may be seconds of data. It they are geographically separated, data loss may be measured in minutes.

Using Microsoft's Volume Shadow-Copy Service (VSS), Neverfail can also roll back databases to a previous known consistent state to recover from data corruption or from erroneously-deleted data.

Monitoring

Driven by the rules for each application that have been previously established, AMF monitors server and storage hardware, networks, and applications:

- Servers are monitored by the Neverfail heartbeat mechanism. The conditions under which a heartbeat failure is declared are determined by the rules established for the application. Rules also specify the action to take upon a heartbeat failure (for instance, restart the application or fail over to the standby server).

- Networks are monitored for connectivity. If access to the public network is lost by the primary server, failover is initiated to the standby server. If the primary server should lose communication with the standby server, further changes are queued and are sent to the standby server upon recovery.

- Applications are monitored according to pre-established rules. Metrics can include queue lengths, response times, CPU utilization, storage utilization, memory utilization, and other indicators of proper operation.

- Rule-based performance is monitored at the user-expectation level. For instance, test transactions might be sent to designated URLs to measure response times.

Continuity

AMF ensures continuity of operations by responding to availability issues determined by its monitoring activity. These continuity actions are governed by the rules and polices established for the application. They typically include restarting the application or failing over the application. If an application must be failed over to its standby, other applications in the ecosystem or the entire ecosystem can be failed over to ensure continued satisfactory performance. User sessions are not lost during a failover – users do not have to relog on. Failover typically takes about two minutes.

Following a failover and the subsequent recovery of the primary server, the applications that had been failed over are switched back to the primary server.

Failover and switchback also can be initiated manually for server upgrades or if a potentially serious condition not covered by the rules occurs. Using manual failover, workload can be shifted seamlessly from one server to another.

WANsmart

Neverfail provides intelligent usage of wide-area networks to minimize the bandwidth required. This not only minimizes replication latency, but it also frees up bandwidth for other uses.

To do this, Neverfail's WANsmart feature has two capabilities:

- It uses data compression to reduce the amount of data that must be sent.
- It uses de-deduplication to further reduce the amount of data to be transmitted. Only the bytes that have changed are sent rather than entire rows or records.

Plug-Ins

Neverfail provides preconfigured plug-in modules to support a wide variety of business-critical Windows applications. Plug-ins provide support for:

- Microsoft Exchange
- Lotus Domino (email)
- SQL Server
- File Servers
- SharePoint
- IIS (Microsoft Internet Information Services)
- Microsoft Mobile
- BlackBerry Enterprise Server (BES)
- RightFax

In addition, support is provided to create rules and policies for custom applications.

Clusters

Neverfail's Cluster Protector provides remote cluster site protection for MSCS (Microsoft Cluster Server) and its successor, WSFC (Windows Server Failover Clustering). Using standard Neverfail continuity services, a cluster can be cloned at a remote site. Neverfail will monitor the health of the cluster and will seamlessly fail over to the remote cluster should the primary cluster fail. Cluster Protector is especially useful to protect a cluster from a catastrophic data-center site failure.

Virtualization

Neverfail can be used between two physical servers (P2P), between a physical primary server and a virtual standby server (P2V), between a virtual primary server and a physical standby server (V2P), or between two virtual servers (V2V).

One common use for Neverfail P2V is to back up several physical servers with a single virtual server. Neverfail will keep each physical primary server cloned in a virtual machine on the virtualized server host. In this way, rather than having one standby server per primary server, only one physical server is needed to back up several other physical servers. Of course, the virtual server host must have the capacity to handle the number of physical servers that may be down at the same time.

Neverfail supports VMware ESX, Citrix XenServer, and Microsoft Hyper-V virtualization environments. Neverfail can migrate virtual machines in these systems to remote standby virtual machines in the event of a host server failure.

In addition, VMware uses Neverfail to protect its vCenter Server, which is the control hub for the vSphere infrastructure.

SCOPE

SCOPE (Server Check, Optimization, and Performance Evaluation) is Neverfail's diagnostic and monitoring utility. It is responsible for performing the comprehensive health checkups that Neverfail uses to determine corrective actions.

During initialization, SCOPE identifies potential configuration and performance issues so that they can be corrected before putting the application into production. SCOPE then monitors the production environment in real time for the same conditions. It discovers changes in configurations and schedules them for replication to the standby server.

SCOPE monitors many different components of the server and application environment, such as memory utilization, CPU utilization, disk utilization, and network traffic. It will issue alerts if it discovers any condition that might result in downtime.

Tertiary Configuration

Tertiary is an extension of the standard Neverfail configuration. It supports three servers – one primary server and two standby servers. One standby server is collocated with the primary server, and one is located remotely. Both standby servers are kept synchronized by Neverfail replication, and both are monitored by Neverfail.

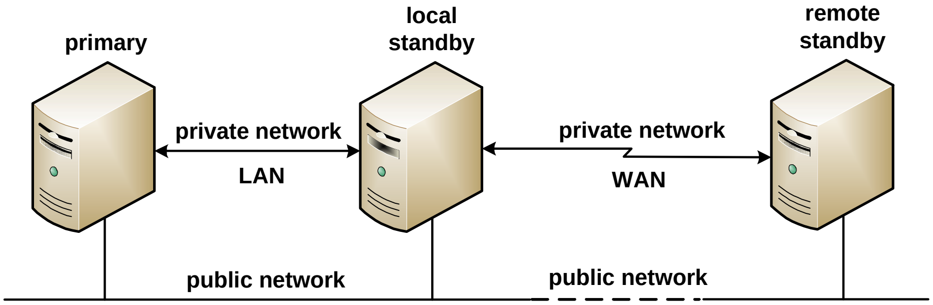

Neverfail Tertiary Configuration

Should the primary server fail, it will fail over to the local standby server if that server is operational. Otherwise, it will fail over to the remote standby server. In this way, failover time and data loss are minimized if the local standby is available. However, the remote standby ensures disaster tolerance should the primary data center suffer an outage.

Continuous Availability Director

The Continuous Availability Director (CAD) is the graphical user interface (GUI) to Neverfail. It provides functions for configuring Neverfail and for monitoring and controlling the organization's ecosystems.

It is through CAD that plug-ins are installed. Rules for component monitoring and failover polices are established for business functions that are not covered by a plug-in, such as special applications used by the enterprise. Neverfail requires no scripting for the rules and the actions that they invoke. All rules are configured via the CAD GUI. For applications that are supported by standard Neverfail plug-ins, no action is required of system administrators since the plug-ins come preconfigured with applicable rules. However, a knowledgeable administrator has the flexibility to monitor and enhance these rules.

Once operational, the ecosystem availability states are displayed graphically so that operations personnel can rapidly note and analyze potential fault issues.

System operators can initiate a variety of functions via the Continuous Availability Director. For instance, if a potential fault appears ominous but has not yet met the criteria for an automated failover, the operator can command a failover. The operator can command a failover if a system is to be taken down for maintenance. The operator can initiate a rollback of a database to a consistent point in time if the database has become contaminated or if a file or table has mistakenly been deleted.

Summary

Neverfail ensures the continuity of critical Windows-based business functions by protecting an entire functional ecosystem. It monitors all components of the ecosystem not only for health but also for performance from the users' viewpoint. If any condition that threatens the continuity of a business function arises, Neverfail will immediately take corrective action, including seamless failover to a standby server that may be locally or remotely located.

Neverfail's monitoring and corrective actions are governed by rules established for each application. These rules are preconfigured for many Windows business-critical applications. Consequently, Neverfail is positioned to protect many critical Windows environments substantially "out-of-the-box."[1]

[1] The material for this article was taken solely from the Neverfail web site. Neverfail should be consulted to confirm the accuracy and the current availability of features described herein.