

Windows 2000 magazin

www.win2000mag.de

Neverfail - The Server's Heartbeat

Server software for reliability

by Johann Baumeister

One of the key features of a server system is its permanent availability. To achieve this, functionality such as mirroring, clustering and fault tolerance is required. We've tested a software solution, which claims to be able to perform these tasks in realtime.

Since server systems are commonly used as a communication node for a large number of clients, their availability is of critical importance. Many different solutions have been developed with the goal of ensuring high availability. The most obvious is to replicate all the hardware components in a box, providing automatic failover to the unaffected component. This method is quite efficient, but at the same time in most cases very expensive, because special hardware is required.

The vast majority of clustering systems use the method of duplicating identical hardware over the network. In order to be able to seamlessly switch over to the secondary system in the event of a failure, you often have to purchase identical devices, and moreover, have them certified. Between these two extremes – cluster systems with replicated hardware and failover as a hardware function, and resuming operations by manually restoring the da-

ta – there are those techniques which provide an ongoing replication of data to a secondary system. In most cases they react with minimal delay, and in the event of a failure they can resume the operation of services in a few seconds. One of these products, that we've tested, is Neverfail Heartbeat.

The basics: Straight forward architecture with two computer systems

The solution has a straight forward architecture: It consists of two computer systems, which complement each other in the event of a failure. The „Primary“ server is the primary active system, whereas the „Secondary“ server functions as a standby. If the active server fails, in a few seconds time the secondary system takes over the functions of the Primary. To be able to do this, the secondary system needs a copy of the operating system with its services and the appropriate applicati-

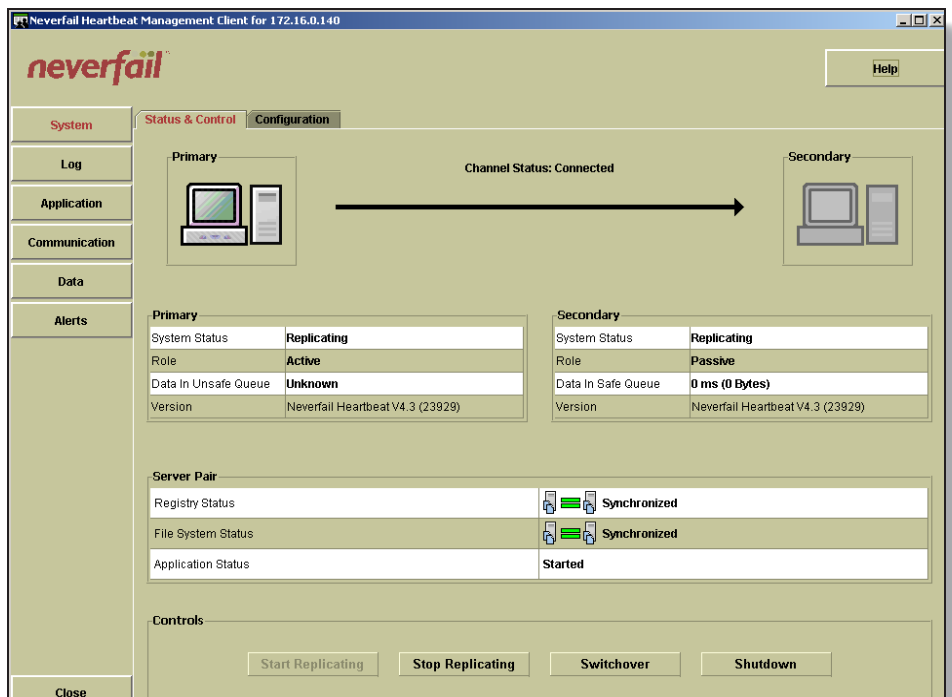
ons. Taking into consideration that in most cases these hardly change, the standby server normally has to be configured only once. On the other hand, the data on the primary server is changing permanently. Thus data is being constantly mirrored to the secondary. Drivers, hooked to the I/O system take care of constantly updating the backup system.

In contrast to other clustered solutions, which our testers have reviewed, Neverfail is quite lenient concerning the hardware requirements. For instance the two systems don't necessarily have to be completely identical. In our test we used systems with differences regarding the types of CPUs, motherboards, disk size as well as network and graphics cards. Needless to say, the differences between the systems have to be taken into account during the installation. However, the basic requirements consist of 512 MByte RAM, two network cards and an identical HAL (Hardware Abstraction Layer) with regard to the CPU of the Windows systems. These moderate requirements allow for even „older“ hardware to be used with Neverfail.

As already mentioned, services and data have to be available in order to supply the uninterrupted resumption of a service on another system. A service in this context is not just an operating system service but can consist of a complete application. In order to resume the seamless operation of a service on the other system, tight integration must exist between the controlling failover functions of „Neverfail Heartbeat“ and the applications. This kind of software integration (the vendor's term is „Application Modules“) is currently available for the file services for Windows, Exchange Server, SQL Server, Sharepoint Portal Server, for IIS-services and BlackBerry's Messaging Server. Additional

functional components (Application Module Extensions = AMX) extend the basic modules and can for example be used to integrate an Exchange virus scanner. We tested the current version, which we downloaded from the website of the distributor Prosoft. The vendor provides comprehensive documentation and a variety of tool sets for checking system requirements. Beyond that, the check lists also serve as useful documentation. The software requirements include Microsoft Windows Server / Advanced Server 2000 with SP4 or Windows Server 2003 and 2 GByte of hard disk space for the program plus the temporary Neverfail data. The solution's Exchange version naturally requires an Active Directory and the DNS service. In addition, the Exchange server can not be used as the domain controller at the same time. Therefore the domain controller together with the „Global Catalog“ has to be installed on a separate system. The solution supports Microsoft Exchange Standard/Enterprise 2000 with SP3 as well as Exchange Server 2003. The SQL Server version also requires the Standard or Enterprise Version 2000 with SP3.

In addition, Neverfail recommend using a system with separate hard disks for the operating system swap files and for the application data or log files. Although this is not a hard requirement with Never-



Ill 1: The admin console: The control panel provides an up-to-date overview of the state of replication.

Neverfail Heartbeat

Vendor:

The Neverfail Group
<http://www.neverfailgroup.com>
 +44 870 777 1500
uptime@neverfailgroup.com

Pricing:

Neverfail Heartbeat Basic Product for one pair of servers with 1 CPU each: 4176 Euro including 12 months support and update service.

Additional software: Application Module for File Server, IIS, SQL, Exchange starting at 1044 Euro per pair of servers
 Extra charges for deployment in systems with more than 1 CPU.

Pro:

- Quick deployment
- High operational security
- Tolerant and flexible hardware requirements
- Comprehensive and useful documentation

Con:

- Deployment limited to Windows systems
- No integration with system management tools

fail, it does help with the performance requirements for the server systems. Neverfail can also be used for the protection of server systems in WAN-environments (Wide Area Network). Thus it is possible to build geographically distributed high availability systems. For that kind of scenario, the vendor recommends the use of a third network card for ensuring high availability even in the case of failures of the network connection in the WAN.

Because the Secondary Server is a clone of the Primary, the software requirements for the Secondary are identical with those of the Primary: They include the same operating system with the appropriate services, service packs, patches, and applications. Moreover, the letters for naming the disk drives, the directory structure, as well as the administration accounts and passwords have to be the same on both systems. The actual installation of the product starts with the „Primary Server“ and is fully menu driven. The software comes with comprehensive documentation and is supplemented by an online demo showing each step of the installation and configuration of the entire system in over 180 animated screenshots.

The configuration: „Heartbeat“ of the systems

When installing the procedures and integrating the server's guest operating system, you have to differentiate between

the Primary and the Secondary. Afterwards some configuration windows open up. One of the most important steps here is the configuration of the „Heartbeat“ between the devices, which is an involved process. Each system must have at least two network cards (NIC), where one set is needed to create a channel for internal communication between the Primary and the Secondary. The user can decide on the port number to be used. The second pair represents the network interface to the public LAN for communication between the clients and the server. When configuring the IP channel, the address of the Secondary Server has to be specified and confirmed by a return. Only then can the installation procedure be resumed.

At this stage the Secondary must be available (e.g. reachable by ping). In our test we were stuck at this point and we broke off the installation for the time being. A call to the distributor's support center showed that we hadn't been the only ones to misinterpret this step. The vendor also supplies an Uninstall application, which returns the system to its original state. So, due to our mistake we unintentionally tested this option as well. The software cleaned up the systems, so that no residues were left behind. This was the only incident needing external support. The comprehensive and useful manual with animated screenshots gives detailed information on each step.

The services, directories, and hence the data to be mirrored on the secondary system must be specified for the configuration procedure. In order to enable a seamless failover to the backup server, it must have a near-identical environment to the primary. This requirement includes primarily the data, regardless of whether it is stored in a file or in a buffer of the operating system. The next step is the configuration, registry entries and of course the applications with their services. The appropriate data files for replication are automatically selected, de-

block file is transferred to the Secondary system. Finally two drivers for the network cards must be configured and then the installation is complete. The user receives a tool icon with the main functions on the lower right hand side of the system task bar.

Next, the Secondary Server has to be configured. Initially, a restore is performed, including the system configuration of the above mentioned data of the primary system. This is done with the help of the Windows restore assistant. After the subsequent reboot, the Secondary becomes a clone of the Primary and has the same software state, system name as well as the same IP addresses. If there are hardware differences between the two, such as a graphics card, the Windows assistant adapts them to the given situation. Finally, as with the Primary, the drivers for the network cards have to be configured for Neverfail.

A further configuration procedure for the Secondary includes adapting the network connection. Due to the setup process, the systems have different channel IP addresses, but the same address for the public LAN and the same system name. In order to avoid duplicate IP addresses or conflicting name resolution, both the DNS registration and the name resolution for NetBIOS over TCP/IP have to be disabled for the Secondary. During changes in the network configuration the system should not be connected to the network. After the changes, a reboot is required and then the system is up and running showing its operation in an icon of the task bar.

The Primary Server is responsible for all further management tasks and for monitoring. The configuration and control of these tasks is done with the help of different functions and filters. For checking the replication between the devices, the admin can start the task manager of the operating system and choose network. (ill. 2). In our test the data replication between the Primary and the Secondary worked fine.

In the next stage, we simulated a failure by disabling the primary server. As a trigger for a failover, Neverfail monitors both the connection to the clients in the network and the functionality of the replication channel to the Secondary. The admin can define a threshold for both possibilities. If one of the two connections fails, a switchover to the Secondary occurs. This protection implicitly includes hardware failures such as those concerning the power supply, the motherbo-

ards or network functions. – because the exchange of heartbeat between the components stops. The failover threshold should be tested very carefully and set to an optimum level, otherwise a temporary bottleneck on the Primary or on the network could lead to an unnecessary switchover. There is a Neverfail tool named SCOPE available for performance assessment, which logs the workload and performance of the system. The data can also be sent to the vendor for analysis. Furthermore, the controlled transfer can be triggered manually.

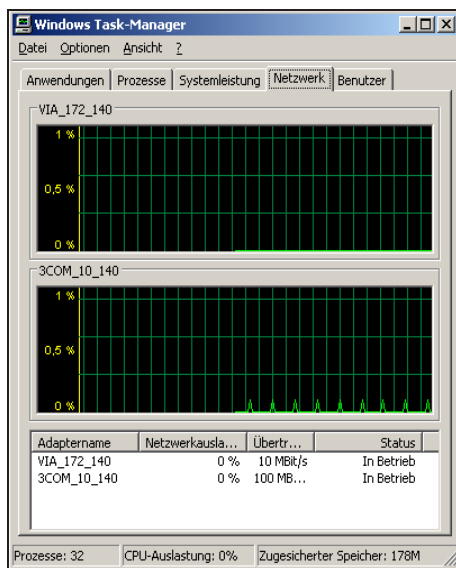
Switchover – the planned takeover

There are several steps in initiating a planned takeover (a switchover) of the active server's role by the Secondary. Firstly, the application on the Primary is stopped, the remaining data is transferred to the Secondary and then the network connection of the primary system is stopped. The Secondary, which is to become the active server, takes over the remaining replicated data, connects to the network and then starts the application.

In the case of a failover, i.e. the complete failure of the Primary server, the steps performed by the Primary are different from the switchover described earlier. However, for the Secondary, which takes on the active role, there are hardly any changes. In order to establish the network connections, Neverfail makes changes to the configuration, e.g. for the DNS name resolution. Starting – and if need be, stopping – the applications is done by scripts, which are created during the configuration.

Summary

During the entire test, working with the product was a pleasant and smooth task. The installation posed – with the exception of the one snag – no problems and took less than half a day plus the time needed for the backup and restore of the data. Despite the clarity and ease of use of the solution the user should still read the manuals carefully and make sure he understands the concepts. The set of tools which comes with the software does its job stably and reliably. With the help of these tools, the product provides the necessary functions for both realtime disk backup and high availability in WANs. Companies can take advantage of the fact that they don't need identical hardware for the two servers and thus have a broad variety of usage scenarios at their hand. (fms)



Ill. 2 The „heartbeat“ over the network: The Windows task manager shows that the heartbeat and the replication are performed through their own network channel.

pending on the version of Neverfail (Application Module: file services, Exchange or SQL Server). For the Exchange version, these are the Information Stores of the mailboxes, whereas for the SQL Server the selected files are the database tables. Any kind of data, including databases or the mail store in Exchange can be replicated. To maintain data integrity in the case of databases, log files must be replicated as well as DAT files. In addition, Neverfail creates six scripts for starting and stopping applications on both servers. Further steps in the configuration procedure consist of specifying the directories, the size of the log files or the public address of the server. Although this is not a really complicated task, one should still follow the documentation. At the end of the installation of the Primary, a backup of all the necessary data is performed, and then the data as a