

# MARATHON EVERRUN 2G PROVIDES SAFETY NET FOR APPLICATIONS HIGH AVAILABILITY FOR THE MASSES

**On the phenomenon of high availability there exist quite a few misconceptions. It is often thought that it is intended to put systems back on their feet after a power failure or unexpected failure. Although recovery certainly plays a role in high availability, it is much more about preventing downtime and data loss. Marathon Technologies shows this with the new everRun 2G.**

In some sectors, such as the banking sector and securities, high availability has already been part of the IT-furniture for a long time. In recent years we see that HA is also increasingly applied outside such demanding environments. Two factors play a role. First, increasing numbers of companies let themselves be carried along in the flow of the 24 hour economy. Certain key issues must be continuously available if a company wants to maintain its position among its competitors. In addition, falling prices and increasing ease of use have a determined pull-effect. HA was until recently focused on companies with large budgets, and then only for company-critical applications. Nowadays more and more companies can afford a HA solution.

## Failed failover

Modern virtualisation techniques are useful for protecting applications from planned downtime. If necessary, you can move a running virtual machine without pain from one physical server to another. But protecting virtualised environments against unplanned downtime is a different story. Traditional approaches such as clustering and failover, fall down in a number of areas.

The main problem is that clustering and failover add a lot of complexity to the environment. To ensure that decisive action is taken in case of an emergency, we have to perform quite a few manual configurations, set-ups and tests. Perhaps not insurmountable, but remember that any additional action increases the chance of errors — and that is just what we are trying to avoid.

Looking at “heartbeat pings”, a widely used control mechanism in cluster and failover

solutions, it turns out that it is not possible all the time to measure reliably the status of a virtual machine. Also, these solutions are not always able to distinguish between I/O- path errors, server errors and a lack of system resources. In some cases such failures lead to unnecessary or even improper failovers. In other cases subtle failures in storage or network devices will go undetected, so that the necessary failover is not executed at all.

Finally, everything depends on a faultless configuration of the systems. If the target system is not configured properly, then after a failover the application or virtual machine will not budge on the standby system. This creates what is sometimes called a “failed failover”.

## Three levels

Marathon Technologies, a supplier of software for high availability and disaster recovery, thinks it can do better. The new product everRun 2G would go beyond the existing cluster and failover solutions. According to the company this software provides high availability for a whole range of company critical applications. Flexibility is paramount in this: for applications such as Microsoft Exchange, SQL Server and SharePoint we can enable fault tolerance at component level, and for specialized applications that absolutely must not fail, we can set up two fully redundant systems that run in “lockstep” mode, so that even when server fails there will be no downtime.

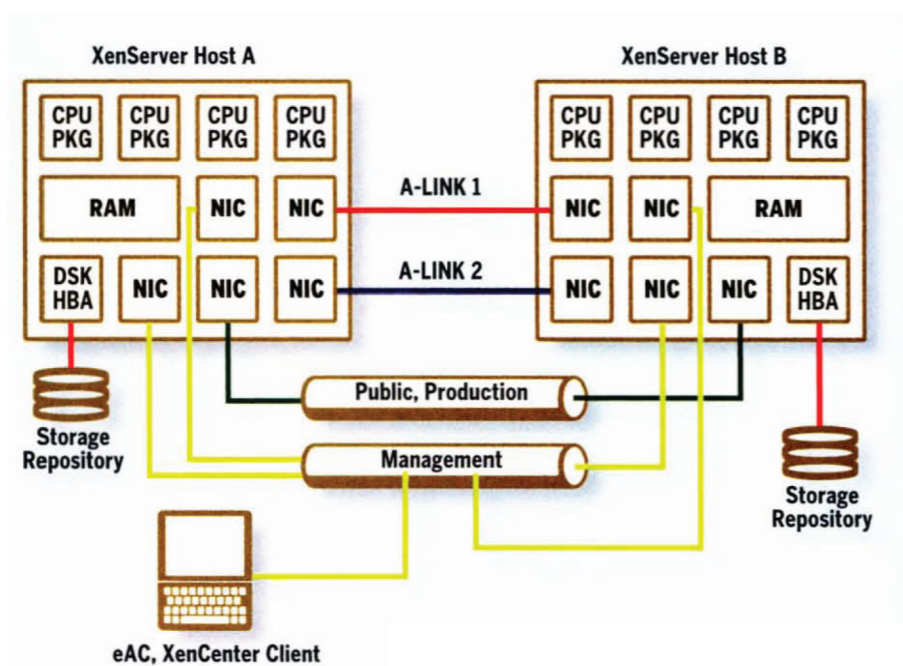


Figure 1. The simplest setup: a local pool with two XenServer hosts.

Marathon everRun 2G was developed in cooperation with Citrix. Concretely this means that everRun uses a Citrix XenServer to provide protection to Windows applications that run inside a virtual machine. This protection consists of three levels. The lowest level, Level 1, is based on Citrix XenServer HA. This level provides basic failover and recovery functionality for applications that are not vital to business operations. Failover takes place to another host in the same pool and there is no storage and data protection. Level 2 provides fault tolerance at component level: applications will survive the break-down of an individual network or storage component without a hitch. There is a synchronous data mirroring between hosts. Unlike Level 1 shared storage is not required. For the real thing, fault tolerance at system level, we have to be at Level 3. This level provides the same type of protection as Level 2, but in the event of a failure the memory status and application status are maintained. Moreover, Marathon promises that even when a complete system crash occurs, there will be no downtime.

## Configuration

Marathon everRun 2G can be configured in three different ways. The first and easiest way is to set up a local pool with two XenServer hosts and a separate workstation running both XenCenter and everRun Availability Center (EAC). The EAC interface allows the administrator to perform various everRun tasks. The two physical servers are directly connected. The first instance of the protected VM is working on one server, the second is on the other server. Two so-called A-links (“availability links”) provide a redundant connection between the two halves of the protected VM. In this configuration, it is not possible to use quorum services, a communications service included in everRun that prevents both halves from operating independently (a situation known as “split brain”). This service is also not necessary because the short point-to-point A-links reduces the chance of a simultaneous network failure to a minimum. The second way is to set up a local pool with two to four XenServer hosts within a single site, with active A-links and quorum services. For this setup, two separate quorum servers are required at strategic positions in the production LAN.

## THINK FIRST, THEN ACT

Prior to the question of what type of HA solution best fits a company, we first have to answer two other important questions. First, what services, data and applications must be continuously available? And, not unimportant either, what does it cost us if we do not have these things available? The requirements are generally expressed in two factors: the Recovery Time Objectivity (RTO) and Recovery Point Objective (RPO). On this basis we can then draw up an SLA. The often used ITIL-KPIs (Key Performance Indicators) for Availability Management are a good tool for testing a SLA. With this information in your pocket, the company can examine which system or product best meets the SLA requirements.

The third way is to set up a so-called Split Site pool with two XenServer hosts in separate data centres or buildings that are geographically separated and connected via a LAN. Again, using quorum services.

Before we start installing, we need to know what level everRun our hardware supports. To that end, we can download the everRun Compatibility Check utility from the Marathon website. This utility checks things like brand and type of CPU. At the end of the check a summary report tells us whether everRun Level 2 and Level 3 are supported.

As a pilot, we set up a pool with XenServer hosts (see Figure 1). As a storage system for both servers, we use the local disks, but it is also possible of course to opt for shared storage on the basis of NFS, iSCSI or SAN fabric. After configuration of the network and storage system, we can begin the installation of XenServer on both hosts. Subsequently we install XenCenter on a server or workstation and include both hosts in a resource pool.

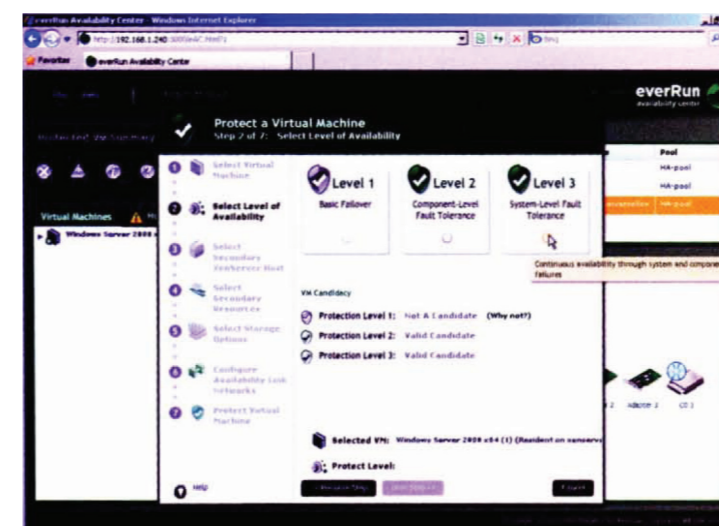


Figure 2. Marathon everRun 2G offers a choice of three different protection levels.

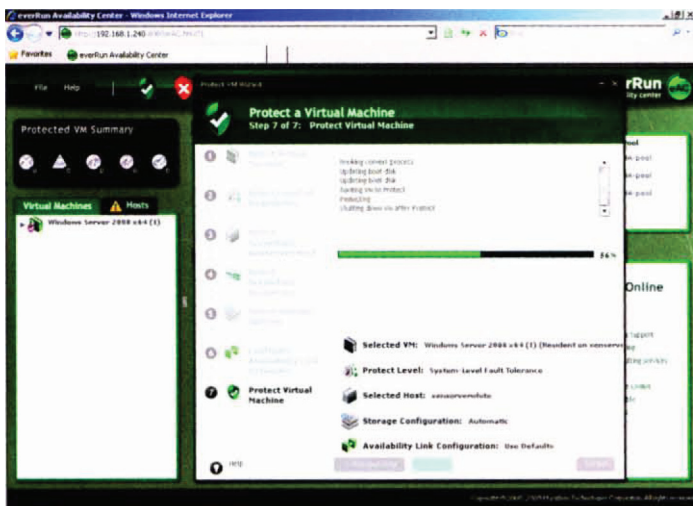


Figure 3. After completing a few steps, the selected VM is protected from interference.

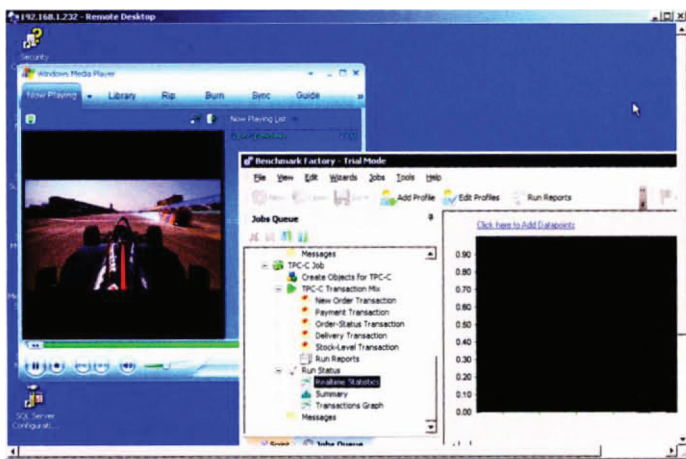


Figure 4. Here we do a little test with two active applications on the primary XenServer host.

Protection for the VMs we created is arranged using the Protect VM Wizard. A matter of performing a few steps and you're done. We opt for Level 3: fault tolerance at system-level, zero downtime for mission-critical applications (see Figure 2). First we select the VM and the required level of protection and then the secondary host that contains the redundant resources. The following menus successively configure the resources on the secondary host, designates the storage resources and provides the A-links with static IP addresses (see image 3). Moreover, in addition to fault tolerance, everRun supports two types of migration to protect the VM: Migration with a reboot and online (live) migration. In other words, the administrator can exit a protected VM and migrate to another host or it can migrate without shut-down.

## Big question is of course whether a VM continues to run without any noticeable effect if one or more hardware components fail

### Test run

Big question is of course whether a protected virtual machine continues to run without noticeable effect, if one or more hardware components fail or even the entire server. Therefore we do a little test on the basis of two active Windows applications on the primary XenServer host (see figure 4). The first application concerns a video running in Windows Media Player. At the same time we performed a TPC-C benchmark on a Microsoft SQL Server 2005 database with the Benchmark Factory package from Quest Software. In this benchmark multiple users simultaneous made changes in the database. Both applications are a quite challenge for most servers: streaming

I/O and a simultaneous application of arbitrary database operations.

Using Remote Desktop on Windows Server, we connect to the Windows Server 2003-VM that is running on XenServer. The aim is not only to verify that both applications actually remain available, but also to see if the remote connection is transferred to the secondary server without problems in case of failure of the primary server.

In the first test, we disconnect in turn various network connections to the primary server, both A-links, the production and management connection and finally all connections simultaneously. We find that in all cases, the remote connection remains intact and both applications continue to run without noticeable effects. In the second test, we disable the storage system on the primary XenServer by disconnecting the local disk from the power supply. Again, no appreciable effect on the running applications and the remote connection. We then start the disk again and reboot the server. After a while the secondary XenServer comes back online and the disks of both servers are synchronized, after which the VM is protected again at Level 3. During the outage we do unfortunately see that the current status of the servers and VM protection at that time is not properly displayed in the everRun Availability Center. In the last test we disable the primary server by disconnecting the system in one fell swoop from the power supply. Even now, the

applications continue to run quietly. Then we enable the primary server again. After some time the server is taken back into the pool and both disks are synchronized.

### Conclusion

In all evaluation tests we have conducted to date, Marathon's everRun 2G is the only software product that manages to provide complete continuity in all circumstances. Failure of system components, power interruptions, the product does not fail. If you wants to try it yourself: an evaluation version of Marathon everRun 2G can be downloaded via [www.marthontechologies.com/high\\_availability\\_software\\_everRun2G.htm](http://www.marthontechologies.com/high_availability_software_everRun2G.htm).