



Being Mindful of Stateful Recovery

By Dan Kusnetzky, Principal Analyst
Sponsored by Scalent Systems

Organizations have long faced the challenges of developing a disaster recovery (DR) plan for their mainframe and midrange computers. Industry standard systems, on the other hand, were not always considered in that plan because they were often performing less critical tasks.

As more and more critical tasks are hosted on industry standard systems, these DR plans must be reconsidered. Furthermore, planning for a disaster isn't enough. These DR plans must be routinely and fully tested. The Kusnetzky Group would strongly suggest that organizations have a well thought out, tested disaster recovery plan **before** a disaster occurs.

DISASTER RECOVERY — THE STATEFUL AND THE STATELESS

IT planners have to consider quite a number of things during the creation of their DR plans including the following: Monitoring the power suppliers, air conditioning equipment, systems, storage and networks to determine if the situation is normal or if problems are appearing. If problems are emerging, the next step is to gather information on applications, systems, storage and the network to determine whether they are they up and running as normal, in the process of automatically recovering from an outage or not running at all. From this point forward, the complexity of the task escalates rapidly.

MODERN APPLICATIONS ARE COMPLEX

In today's world, application systems are seldom monolithic processes that run on a single industry standard system. Most modern applications are constructed of a number of tiers, sometimes called "services", such as, presentation services, application processing, data management and storage management. While some of these functions may reside on a single system, it is much more likely that each of them is hosted separately.

Furthermore, to gain increased scalability and reliability, applications are increasingly being segregated into those that are stateful and those that are stateless.

WHAT DO THE TERMS "STATELESS" AND "STATEFUL" MEAN?

Stateful and stateless are adjectives that describe whether a function is designed to note and remember the results of one or more preceding functions or events in a given sequence of interactions with an individual, another function or, perhaps, another system. Stateful means the function is designed to keep track of the state of interaction, usually by setting values in a storage field designated for that purpose. Stateless means that the function is designed so that it doesn't need to keep track of outside events.

As a the need for a given application or workload increases, stateless functions may be replicated, that is multiple instances started up in the network, and workload management or orchestration software can balance the use of each of these instances to increase overall scalability. This has a wonderful side effect of

making that function much more reliable as well. If one instance of a function or application stops functioning due to a planned or unplanned outage, individuals using the function or application are able to continue their work. The workload management system simply forwards incoming requests to remaining instances of that function or application.

Increasingly these functions or, perhaps, the entire stack of software comprising the application system, may be encapsulated into a virtual machine.

Creating a workable datacenter disaster recovery entails awareness of recovering stateful and stateless functions on each tier of a complex application.

ORCHESTRATION OF STATELESS ASSETS

The first step is isolating the stateless processes from the stateful processes so that each can be handled appropriately. This means repurposing servers, reconfiguring network and storage services in real time.

MONITORING

To accomplish this feat, IT executives know that they must adopt technology that allows application components as well as the underlying physical and virtual systems to be carefully monitored. Monitoring only the health of the physical systems that support multiple application components, applications and/or virtual systems just isn't enough.

OPTIMIZATION

It is not enough to gather the details concerning the state of all of the application components, applications, virtual systems and physical systems. There is just too much information for the IT administrative staff to monitor in real time. This information must be integrated, decisions must be made on the appropriate actions and these decisions must be put into immediate action. Unfortunately no human being or group of human beings knows enough about what's happening inside of a complex computing solution or can do this fast enough to make any necessary changes to the environment before an application slow down or failure is seen by those accessing the solution. Tools are needed that can integrate the information about all of organization's applications, application components and virtualized resources and then make the appropriate decisions enabling these applications to always meet services level objectives and perform according to the organization's policies.

AUTOMATION

Making these key decisions isn't enough. People simply can not act fast enough unaided. So, other tools must be deployed to act based upon the decisions made by optimization technology. This means giving high priority tasks more resources (processing time, memory, storage and the like) when it appears that their performance is not going to meet the minimum requirements for that application system. This also means reducing the resources allocated to lower priority tasks when that is necessary.

Real time adjustment of resource assignments must occur without requiring a great deal of staff time, attention or expertise.

SERVER REPURPOSING FACILITATES DISASTER RECOVERY

Since it is possible to provision a physical server from the bare metal, load the appropriate software, configure it to use the appropriate storage and place it in the network in five minutes or less using the appropriate combination of virtual processing software and management software for virtual environments, The Kusnetzky Group suggests that this approach be a central part of IT disaster recovery plans. Once in place, this technology would make it rather



straightforward to resolve outages using an automated, policy and priority driven plan.

Decision-makers would be well advised to include server repurposing and its requirements in their planning regardless of whether the organization is planning to deploy Microsoft's, VMware's or XenSource's virtualization technology. Scalent Systems, an industry leader in server repurposing technology, offers the tools to make server repurposing straightforward and dynamic. More information on Scalent System's approach can be found here: <http://www.scalent.com/kusnetzky>