



Automated Disaster Recovery for KVM-Based Clouds


As more organizations migrate from VMware to KVM, they're discovering that VMware provides some solutions that are absent in the KVM ecosystem. One of the biggest gaps is that there is no KVM-equivalent for their disaster recovery solution, VMware Live Site Recovery (VLSR), formerly known as VMware Site Recovery Manager (SRM).

The problem is that Linux KVM has no cross-site intelligence – there is no single point of management requiring each site to be managed separately. The hosts, clusters, virtual machines (VMs), and clouds have no knowledge of anything running in another data center. Without a built-in Disaster Recovery (DR) solution, organizations are left with two options, purchasing and installing a third-party solution, or taking a do-it-yourself (DIY) approach.

The Disaster Recovery Engine is Built-Into StorPool's Software-Defined Data Storage



With the StorPool Disaster Recovery Engine (DRE) – a built-in component of the StorPool Storage solution, organizations can configure policies for data replication to remote sites and automate failover (and failback) between sites simplifying the development of a comprehensive DRaaS offering managed through the service portal of a service provider or their CMP.

DRE Deployment Options




Active-Passive - An Active-Passive solution is the classic DR example, using a primary site and a recovery site. During a failover, operations at the primary site would be stopped and the recovery site would use the VM recovery points to start everything up at the recovery site.

Bi-Directional - A Bi-Directional solution was the first evolution of classic DR. In this model, customers would run two primary data centers, each running different workloads. Data replication would be running in both directions, allowing each site to act as the recovery site for the other. During a failover, the workloads would both be running in one of the data centers until the disaster was remediated and one of the workloads would fail back to its original site.



Many-to-One - A Many-to-One solution is often chosen as a cost-effective option. In this model, a single site would act as the recovery site for two or more primary sites. This allows customers to have only one recovery site instead of a one-to-one data center configuration. Since most disasters happen only within a specific local region, the recovery site would only require enough compute resources to act as the recovery site for one of the primary sites, not all of them, reducing the cost of the overall DR project. This DRE option is not available in most DR solutions.

Multi-Site Mesh - A Multi-Site Mesh solution allows for load-balancing across multiple recovery sites. You can think of it as a combination of multiple pairs of Bi-Directional sites configured into a larger solution. DRE is very configurable. For example, in the diagram shown, one of the primary sites, running 3,000 VMs could be configured to replicate and failover 1,000 VMs to each of the other three sites. This model can be ideal for customers running and supporting multi-region clouds. This DRE option isn't available in other DR solutions.






We have trusted StorPool for over a decade to provide key storage components of our clouds worldwide. StorPool fuses operational expertise with a market leading product for price-performance. With the new Disaster Recovery Engine, StorPool is helping us to better meet the needs of our customers with improved consistency of service.

Robert Jenkin | CEO with CloudSigma

Disaster Recovery Engine In Action

Disaster Recovery Engine (DRE) makes it simple to create and enforce data replication policies, automatically creating VM recovery points, as well as automate failover and failback between remote sites. Currently, DRE automates failover – making it a turn-key DR solution. Automated is different from automatic. An automatic solution would cause a failover to the recovery site(s) when it detects a data center outage, while an automated solution requires that a human being make the explicit choice to initiate failover and failback action. Someone working in the data center will be able to discern between a momentary outage, like a glitch in electric power, that will be over in a minute and a full-grid electric power failure where it might be days before power is restored.



No Additional Licenses Needed

StorPool Disaster Recovery Engine is built into the StorPool Storage software and is included as part of our basic license at no additional cost.

The only caveat is that StorPool does need to be installed at both the primary and recovery sites, and the same cloud management platform needs to run on both systems. Each management platform has its own format for virtual machine metadata. Since DRE policies can be set on a per-VM basis, it is possible to be running two clouds at each site, each with its own cloud management suite; the VM restore points just need to be activated in the same cloud platform at the recovery site.



www.storpool.com



+1 415 539 0356
+359 (0) 88 834 6465



info@storpool.com