

Tintri Synchronous Replication

Industry-Leading Functionality Puts Zero RPO Replication Within Reach for Every Organization's Mission Critical Applications

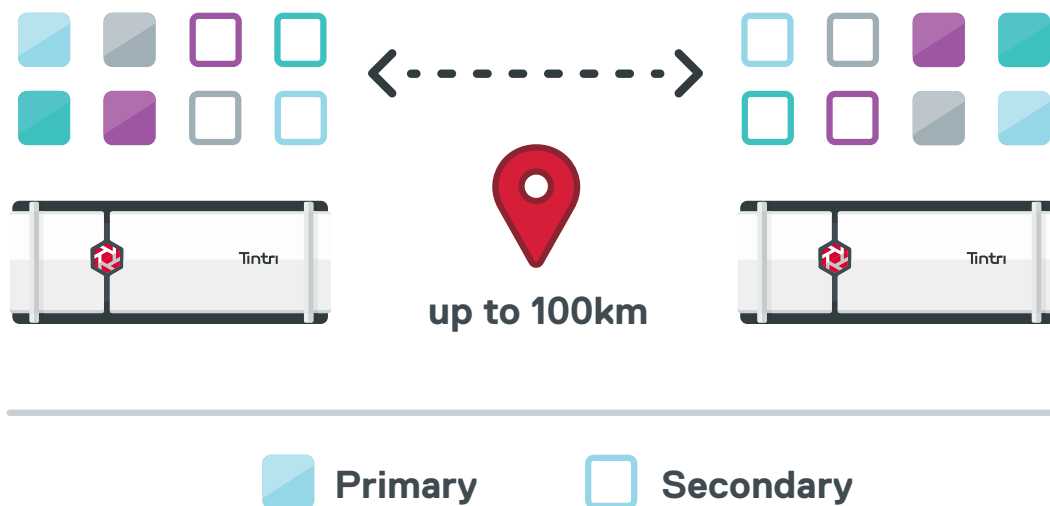
Synchronous Replication is essential to the recovery of mission critical virtualized and cloud applications. Until now, organizations that wanted synchronous replication to protect those applications had to use complex and expensive storage, limiting the benefits of zero recovery point objective (RPO) disaster recovery (DR) to those with the largest IT staffs and budgets. Tintri Synchronous Replication provides those benefits with simple set up, one-click failover, plus all the operational and performance advantages of Tintri VM-aware Storage (VAS) for virtualized and cloud infrastructures.

Zero Data Loss DR for Mission Critical Applications

With Tintri, synchronous replication can be set up in minutes by an IT generalist and offers unmatched:

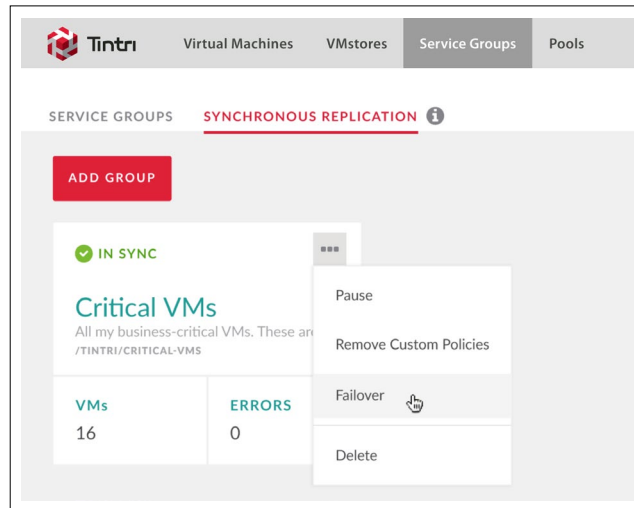
Flexibility

- Recovery Time Objectives (RTO) as low as 30 seconds or less for planned and manual failovers
- Support for inter-data center distances up to 100 km or round-trip time of 10ms across a metro cluster network
- Support for synchronous and asynchronous replication on the same array simultaneously
- Bi-directional synchronous replication, where an array can host primary VMs and replica (secondary) VMs for different groups of VMs simultaneously. In other words, a VMstore can be the primary and replication target at the same time.



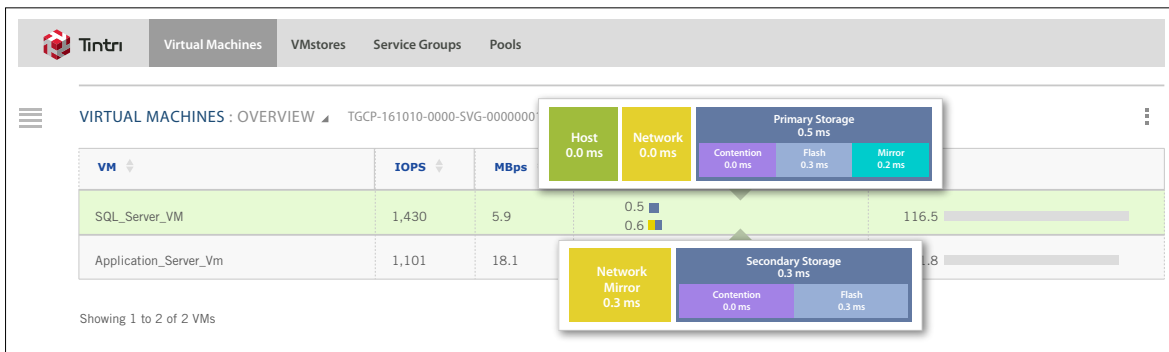
Simplicity

- Does not require purchase and configuration of additional hardware such as FCIP converters or SAS switches, or specialized array models
- Easy configuration and single-click failover enable IT generalists to manage disaster recovery and data protection for critical applications. All you need to do is right-click on “Failover” for the VM group.



Visibility

- Real-time analytics for synchronous replication at the array and/or VM level with end-to-end visualization of latency—including the replication component. See figure below for latency breakdown of both primary and secondary storage, along with host and network.
- Monitoring statistics at a granular level such that the percentage of time that the synchronously replicated datasets are “in sync” can be clearly displayed. This is crucial to capturing “out of sync” incidents quickly and accurately.



“Tintri has architected synchronous replication to not only bypass the single point of failure, but it also allows you to have an RPO of zero. If you find yourself in a situation where you want to failover your storage, you don’t want to be searching around for commands or get lost in a complicated GUI. Tintri addresses this with a single click. Tintri’s synchronous replication is yet another example of how to keep storage simple.”

Matt Crape, C3 Group

