

WHITE PAPER

Enabling **Always On** Business Operations with Guaranteed **100% Data Availability**

The Challenge

As data becomes not only a business enabler but also the base of the business's competitive advantage, Always On data availability has become a critical business requirement. With multiple potential points of failure, IT organizations have struggled for years to minimize the cost and complexity of designing data infrastructure that are Always On. To be truly Always On, a solution must provide:

- ▶ Non-disruptive upgrades, repairs, and refreshes
- ▶ Zero Recovery Point Objective (RPO=0)
- ▶ Zero Recovery Time Objective (RTO=0)

These are necessary requirements from an Always On solution but by no means sufficient. Additional considerations are:

THE PHYSICAL FAULT DOMAIN

Storage systems are always limited by their physical location - if a critical data center infrastructure fails (e.g. power, network, WAN) the storage array may still have the data protected but not accessible to the hosts. Any solution tackling this exposure must be geographically spread to handle the possibility of a datacenter-level failure.

PERFORMANCE OVERHEAD

Any solution spread across geographies with a zero RPO will add some performance overhead as data needs to traverse the WAN, which adds the bulk of latency to the solution. If the solution treats dataset as "local" to one site and not the other, it will further increase the performance overhead when working on datasets from the remote site as data will traverse the WAN twice, almost doubling the latency. Any solution for zero RPO needs to be truly Always On with minimal performance overhead to allow more applications to benefit from automated failover.

AFFECTED APPLICATIONS

With high performance overhead (latency) dramatically affecting many applications due to parallelism limitations (databases especially), Always On solutions are often limited in use for tier-1 applications that are critical to maintaining business operations. Leaving all recovery of tier-2 applications, which often are important to the customer experience, to manual or semi-automated recovery. A solution must have such a small overhead that it can benefit tier-2 and even tier-3 applications, better protecting the business by minimizing impact during failures.

COST OVERHEAD

Beyond the cost of storing 2 copies of the data, there are direct costs to protecting the data with an Always On solution: the need for a dedicated product (HA gateway) or the need for a Fibre-Channel (FC) fabric between the sites to provide low latency. Procuring the HA gateway (often licensed by capacity) and the ongoing cost of the FC infrastructure add friction to adopting Always On solutions. The solution should rely on cost effective Internet Protocol (IP) infrastructure and be fully integrated into the storage layer, without capacity-based licensing, if the goal is to minimizing cost and enable wider adoption.

ADMINISTRATIVE OVERHEAD

When an Always On solution requires a dedicated gateway, it adds complexity to administrative processes:

- ▶ Moving an existing LUN into Always On infrastructure often requires a complete data migration and configuration change.
- ▶ Managing another layer in the application stack, including licensing, monitoring, etc.
- ▶ Complexity of migrating off an HA gateway
- ▶ Data management tasks (e.g. backup & recovery) are performed in the HA gateway while capacity management tasks happen in the underlying storage. This creates inefficient administration.

Any Always On solution should be fully integrated into the storage layer to minimize the administrative overhead.

LIMITED BUILDING-BLOCKS

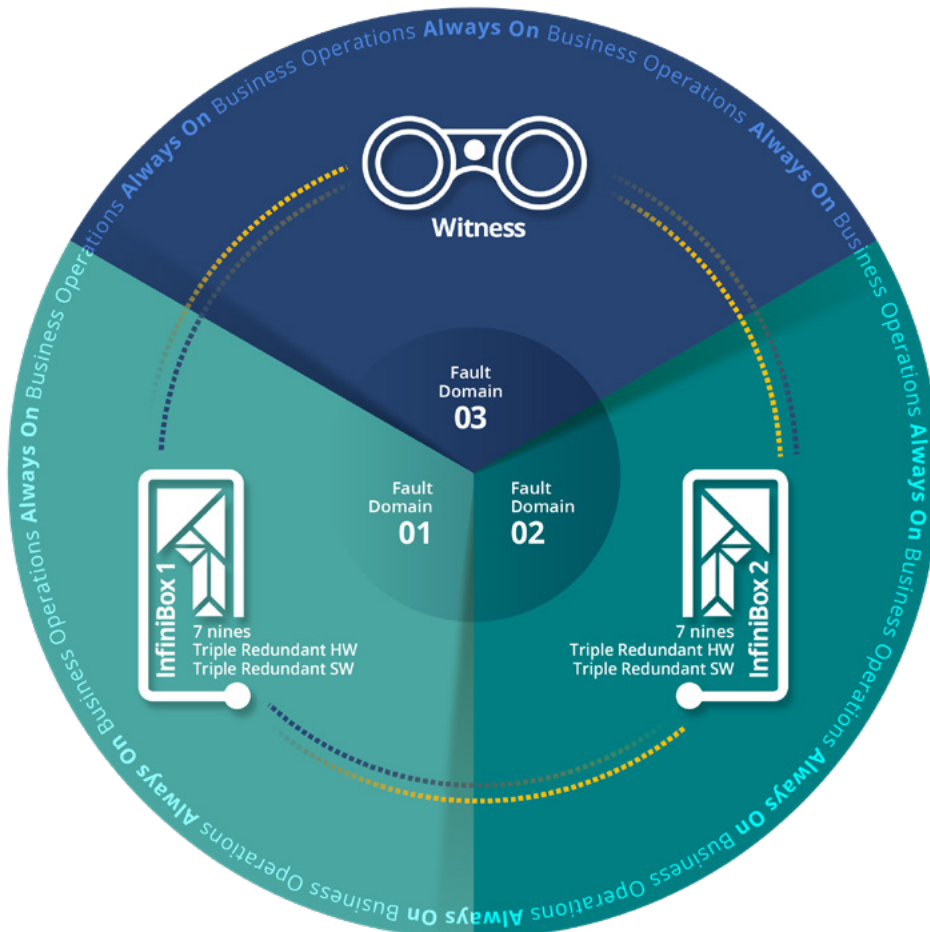
One of the biggest challenges in building an Always On data infrastructure is the building blocks themselves — most on-premises storage arrays are still dual-controller architectures whose lineage dates back to the '80s and '90s when storage arrays were designed for a few Terabytes (TB) and haven't changed much since. Putting hundreds of Terabytes (TB) on these architectures increases the exposure to data unavailability.

The Solution

Infinidat's InfiniBox architecture has been designed from the ground up as a triple-redundant symmetric-access solution, doing away with the old dual-controller architecture and providing customers 100% data availability.

The InfiniBox Always On solution answers all the requirements listed above, and comprises of 2 InfiniBox systems plus a witness spread across 3 separate fault-domains with redundant IP networking between them. This enables customers to a move off traditional dual-controller architecture to triple-controller in each site with another layer of redundancy above them. This architecture resulting in the highest level of availability as failover between relies on simple path failover — just like an HBA / switch failure!

Always On is organically developed as part of InfiniBox, doesn't require a separate gateway or a license. This solves all of the challenges above, while at the same time providing the lowest performance overhead for Always On in the industry.



Building an Always On Infrastructure

InfiniBox's triple-redundant hardware design and 100% data availability as detailed in the [Infinidat Storage Architecture white paper](#) provides the highest level protection within the same fault domain and is augmented by the Always On solution for additional resiliency:

FAULT ISOLATION BETWEEN INFINIBOX SYSTEMS

Each InfiniBox acts as a self-contained cluster, without any dependency on the other box to prevent exposure to any form of rolling disasters. For example, this architecture allows one system to lose power (and consequently all network access) and its Always On partner will not need to wait for clustered services to failover — it can continue operations within seconds.

WITNESS

The witness is provided as a simple to deploy OVF (Open Virtualization Format), with minimal configuration. The witness can be deployed in a customer site or in the cloud. The witness provides protection from split-brain scenarios, helping the InfiniBox system react gracefully when the redundant network between them is down.

PERFORMANCE

In order to minimize the performance impact of the Always On solution, InfiniBox leverages both a direct connection to the optimal remote InfiniBox controller (to avoid east-west traffic) and unique pre-processing of the data that minimizes the response time of the remote system, and avoids most of the performance penalty of an Always On solution.

The Guarantee: Always On with 100% Data Availability

Infinidat offers its customers the Always On guarantee with 100% data availability for data protected by Always On, in accordance with the [terms & conditions](#).

Value to the business:

- ✓ Improved service to customers due to higher application uptime
- ✓ Completely automated Business Continuity using HA instead of failover (e.g. SRM)
- ✓ Triple-redundant array in each site, Fault Isolation between the 2 arrays
- ✓ Automated failover with split-brain protection using witness in third site
- ✓ Guaranteed data availability with clear remedy in case SLA is not met — no fine print
- ✓ Aligns customer and Infinidat interests: we both want your applications Always On!