


WHITE PAPER

# Data Protection

A Comprehensive Approach  
to Protecting Data



## Abstract

Data is the lifeblood of business. Protecting that data is arguably the most important task IT has in the data center today. Data protection has many aspects. These aspects include data availability on primary storage, the long-term preservation and retention of data for multiple business purposes, including compliance and regulations, and the rapid recovery of any and all data from cyber threats. The value of data is increasing and there is more of it being generated every second, compounding the issues of data protection. This is why the technology of protecting data has evolved significantly over the past 30 years.

The requirements for protecting data still have the same fundamental aspects. Clients want their data when they want it and where they need it. To deliver on these demands as data scales, some things need to change. This paper reviews Infinidat's innovative software-defined storage architecture, InfuzeOS™, from a data protection perspective. InfuzeOS enables 100% guaranteed availability, space-efficient, incredibly fast snapshot capabilities for quick operational recoveries, high-performance backup recovery, synchronous and asynchronous replication, and guaranteed cyber storage resilience.

## Introduction

The growth of data makes it hard to keep up availability and backup SLAs. Traditional models of backup and incremental restores are not keeping pace with data protection expectations at multi-petabyte scale. Additionally, organizations are being pressured to increase the responsiveness of their data protection process while driving down overall IT costs.

**All data needs the same level of high protection. You never know what data you're going to need, when you're going to need it, or what you may need it for.**

An important consideration is that the dynamics of the data center have become increasingly complicated. There are numerous applications, databases, and file systems that generate different types of workloads. In addition, as new applications come on line, there need to be ways to protect the newly created data without creating a customized backup solution each time. These dynamics have frequently led to the creation of multiple data protection silos with multiple products within the data center leaving gaps in data protection practices, increased complexity, and exorbitant costs. Almost every IT organization struggles with these challenges.

This white paper will discuss many of the capabilities within InfuzeOS, Infinidat's unique and powerful software-defined storage architecture, as these features pertain to delivering data protection. InfuzeOS is deeply ingrained in every aspect of Infinidat's storage platforms resulting in enterprise-class capabilities that are second to none in performance, cyber storage resilience, reliability, and availability. All the components within InfuzeOS work together intelligently and autonomously to deliver exceptional data services capabilities.

## Protection Starts with High Reliability

### 100% Availability

At Infinidat, data protection starts when the data lands in the array. In a storage market with vendors touting 'nines' (99.999%) of availability, Infinidat guarantees 100% data availability for its primary storage platforms. This level of availability is what organizations expect and IT has previously been challenged to provide. The first step in delivering data protection is building your strategy on the strength of Infinidat's innovative InfuzeOS architecture.

### InfiniRAID®

InfuzeOS protects customer information using a dual-parity, declustered "RAID"-like process called InfiniRAID. It employs a unique approach by creating a large number of data "RAID" groups and automatically distributing its members across all the system's enclosures and disk drives. This method is also known as dispersed data layout and optimizes such data distribution in a 14+2 dual-parity RAID 6-like protection. InfiniRAID enables more of every drive in the system to be utilized which means more usable capacity per system, driving down costs.

Traditional clustered RAID protection schemes rebuild a failed disk from parity within the RAID group which is a limited set of disks participating in the rebuild process. A disk rebuild that occurs within one RAID group, creates a bottleneck that impacts the performance and results in a significant drop in latency across the entire system.


**InfiniBox employs a unique approach of creating a large number of data "RAID" groups and automatically distributing its members across all the system's enclosures and disk drives. This means more usable capacity per system.**

InfuzeOS takes a different approach. Contrary to the clustered protection approach, InfiniRAID rebuild spans the required writes across the entire system; there is no hotspot on rebuild which nullifies any performance impact as each drive in the system helps to perform the rebuild. As a result, disk rebuilds happen in a much faster timeframe.

InfuzeOS's declustered protection maximizes the data access of the system under all conditions, and provides fast data re-protection (disk rebuilds) in an average of 15 minutes for two failed multi-terabyte drives. When two disks fail, we first rebuild those declustered "RAID" groups that have data on both disks as quickly as possible, ensuring all the data is accurate and available in case of another disk failure. All remaining disks in the system participate in the rebuild process and hence the extra load per drive is approximately 0.5%.

### Data Integrity Checksum

In order to extend the resiliency and data protection of Infinidat's systems an extended metadata checksum was created that is applied and added to each data block before it is destaged (written) to disk and re-evaluated every time it is read. Each time data is accessed we are constantly checking and verifying the data to ensure the highest degree of data integrity. Each time we write or rewrite data (for example during a disk rebuild), we also recheck the metadata. This further extends data protection beyond that of any other storage array from the time the data lands in the InfiniBox gives our customers confidence in data integrity throughout their data lifecycle.



## Snapshots and Replication

Using InfiniSnap internal snapshots for operational recovery is becoming more and more commonplace. This doesn't mean that data isn't also "backed up" for longer term and off-site retention. However, to meet the necessary backup and recovery windows, snapshots are the speed-enabling technology.

**The InfiniBox system, with its patented innovative snapshot technology, is designed to provide higher availability, not only for active data, but also for snapshot data.**

The InfiniBox platform, with its patented innovative snapshot technology, is designed to provide higher availability, not only for active data, but also for snapshot data. InfuzeOS's snapshot and replication capabilities help provide RPOs, RTOs, VROs (version recovery objectives) and GROs (geographical recovery objectives) that are unmatched in the industry for uptime and data availability.

An important differentiation of Infinidat's snapshot technology is streamlined data management. Other storage architectures spend a great amount of time "locked," not servicing I/O and buffering data in their drivers while doing data management operations in an effort to support data consistency. Infinidat bypasses all of this. We don't lock data structures in memory. Our virtual user address space (VUA) is made up of multiple addresses that can point to the same piece of data. Creating a snapshot is a pure metadata operation. There is no data movement or copying of data. Regardless of capacity, even at petabyte scale, we create a snapshot without any locking of metadata.

Today, InfuzeOS supports 100,000 snapshots (there is no theoretical limit other than local resources) with zero impact to performance. The system can perform as many as 25 snapshots per second and maintain without latency impact. If clients were to try this on any other commercially available system today, they would see a spike in latency—if they were even able to keep up and do the snapshots at all. This is all due to the locking and unlocking mechanisms that ensure data integrity. Organizations can now create new business continuity metrics and unlimited RPO performing all operational recoveries from a snapshot delivering much higher risk mitigation. In addition, mounting a snapshot or replica and running off a clone happens in seconds, creating a near-zero RTO. Infinidat delivers zero penalty snapshots that are highly space-efficient, and we don't charge extra for this feature.

## Highly Efficient Replication

Infinidat built its replication engine on top of our innovative snapshot engine. The unique InfiniBox destage strategy is a full cache operation for replication management, especially for small RPO intervals. The way replication works is when you mirror a volume, InfuzeOS takes a snapshot every n-number of seconds, sends and applies that log to the remote system. Once there is consensus that the snap has been applied, we roll forward and delete the previous snapshot. As a result, clients can have a system with ~250 replicas and as many as 2048 volumes, or one large volume replicating the entire array or data set with a guaranteed 4-second RPO.

**Asynchronous Replication:** Enables 4-second Recovery Point Objective (RPO). Using an IP infrastructure reduces cost and complexity.

**Synchronous Replication:** Enables zero-second RPO with latency below 400 microseconds for mission-critical applications. Should the WAN lag or fail, InfiniBox synchronous replication fails back to asynchronous mode. When the WAN is restored, the engine will automatically replicate all missing data and resume sync replication without disrupting I/O.

**Active-Active Replication:** InfiniBox systems enable simultaneous read and write to consistency groups over metropolitan areas. The volumes are external images that appear as multi-paths to the same volume. Synchronous replication always keeps volumes consistent. There is no primary-secondary relationship, and no extra round trips to perform write updates to any volume. If needed, an external, lightweight witness can exist on a stand-alone node or cloud-based VM. This active-active replication is at the core of our seamless and non-disruptive in-family data migration.

**Concurrent Multi-site Replication:** InfiniBox can simultaneously replicate consistency groups from main replication sites to another site in a metro area. From there, users can asynchronously replicate to a third remote location.



## Infinidat InfiniGuard

Infinidat's InfiniGuard is built on the InfiniBox architecture and InfuzeOS. InfiniGuard is a purpose-built modern data protection and restore platform. InfiniGuard offers unrivaled speed and reliability in data backup and restore capacity for high-performance business continuity and disaster recovery. InfiniGuard's modern data protection solution is enhanced with InfiniSafe® technologies built-in for cyber resilience, providing a fully orchestrated capability creating immutable snapshots, logical air gap, a fenced forensic network, and near-instantaneous recovery capabilities, at no additional cost. InfiniGuard provides a comprehensive modern data protection solution built on our market-leading InfiniBox to optimize performance and target-side dedupe processing, enabling you to maximize effective capacity for multi-petabyte datasets.

InfiniGuard becomes the repository target and high-performance recovery platform for industry leading backup applications such as: Veritas NetBackup, IBM Spectrum Protect, Commvault, Veeam, Oracle RMAN Backup and other backup control applications.

### Backup and Recovery Performance

InfiniGuard provides industry-leading performance for data protection and backup with in-line ingest rates of up to 180TB/hr\* and unrivaled recovery performance at scale, be it a single file, large-scale DR, or cyber recovery. Cyber resilience best practices require you to ensure data validity, as well as near-instantaneous recovery of your backup repositories, regardless of size. Our InfiniSafe technologies do just that, in minutes, making it easy for data validation and DR recovery when a cyber event occurs. InfiniSafe is detailed in the following section.

### Reliability

InfiniGuard's deduplication engines (DDEs) are fully active/active with a ready standby node built in. We leverage a stateless recovery model which makes it quick and easy to ensure your backup and recovery system is always optimized. Remember, this multi-engine architecture leverages Infinidat's internal full data integrity.

## Consolidation

InfiniGuard solutions scale up to 50PB+ \*\* of effective capacity. With multiple protocols supported (VTL, NFS, CIFS, OST, RMAN, and DB2) it is easier than ever to consolidate data protection for multiple systems in a single high-capacity solution. Consolidation with InfiniGuard also reduces data center footprint and lowers power and cooling costs – significantly reducing data protection TCO. All this while delivering industry-leading performance.

## Security

InfiniGuard provides extensive security features that ensure your data is protected at many levels. Role-based access controls (RBAC) and multi-factor authentication (MFA) integrations are just the front end of security that extends through the entire solution. InfiniGuard's architecture builds upon the deep layers of security inherent in our InfiniBox architecture such as at rest encryption, encrypted replication, and isolated storage for the most critical of datasets using immutable snapshots. All of these capabilities are proven in some of the largest data center infrastructures in the world.

## Infinidat InfiniSafe

The Infinidat InfiniSafe architecture represents four critical areas of data security that are fundamental to cyber resilience. Each of these elements are independent of each other and can be implemented individually or collectively to create the desired level of security that aligns with your organization's business protection goals. InfiniSafe technology is available across the Infinidat portfolio of InfiniBox, InfiniBox SSA, and InfiniGuard.

Creating locked and unchangeable copies of your data is of the utmost importance. These can be considered logically air-gapped on their own, but extending that via a replication best practice to a second, remote immutable copy is important, just like recovery data separation in a DR strategy. Then you need to test and/or validate your data in that copy in a protected environment. Having a fenced environment, sometimes referred to as zero trust, separates you from production and is only active during the time needed to validate what you specifically want to make sure is clean. You can use inspection tools and applications that are best for you to validate and or test the data. Finally, once you have validated those points in time, you can recover that data in seconds to minutes. Leveraging our capabilities within our InfiniBox family gives you all of this with no proprietary need or lock-in to a particular vendor or toolset.

The four pillars of InfiniSafe technologies are:



### Immutable Snapshots

This snapshot capability allows the creation of numerous point-in-time (PIT) copies of data in the most efficient manner possible. These snapshots provide absolute security of the data. There is no ability for data to be changed or altered and snapshots can only expire based on the retention policy.



### Near-Instantaneous Recovery

InfiniBox and InfiniGuard use InfiniSnap technology for non-impacting instantaneous access to the locked data within the immutable snapshots. Recovering data from InfiniBox primary storage is achieved in less than a minute and from InfiniGuard secondary storage in less than 20 minutes, guaranteed.



### Fenced Network Forensic Environment

Creating an InfiniBox fenced environment consists of isolation of needed secure server resources that would connect to a private network via FC or Ethernet within InfiniBox. Access to this private network can be dynamically created via API automation or through the InfiniBox GUI. Secure resources can have immediate access to the immutable snapshots.



### Remote Logical Air Gap

InfiniBox creates a local logical air gap. Using InfiniGuard, you can create a remote local air gap. These air gaps provide organizations with the ability to identify datasets that are deemed critical and create a cyber resilient copy strategy for these data sets.

## Guarantees

Infinidat's comprehensive guarantees for the performance, availability, and cyber recoverability of our platforms set the standard for enterprise storage. IT is challenged to deliver more with less in an increasingly aggressive SLA environment. IT deserves more than casual assurances from trusted vendors. Here are the guarantees available from Infinidat.

**100% Availability Guarantee.** No "9's," just 100% continuous uptime.

**Performance Guarantee.** Written guaranteed performance SLA based on a detailed workload analysis.

**Guaranteed recovery from InfiniSafe immutable snapshots.** Coverage on InfiniBox, InfiniBox SSA II, and InfiniGuard platforms.

**InfiniSafe recovery time guarantee.** Recovery from immutable snapshots in less than 20 minutes for InfiniGuard and less than 1 minute for InfiniBox and InfiniBox SSA, regardless of snapshot size.

## Conclusion

Organizations have very complicated data protection environments which cannot be fully covered or satisfied by a single protection strategy. Infinidat is sensitive to these challenges and offers comprehensive data protection features and options to address the complex needs of modern data protection.

Leveraging the strength and innovation of the InfiniBox internal data protection architecture and InFuzeOS capabilities, 100% availability, and storage performance during data protection is guaranteed. Organizations have the flexibility to use both InfiniBox primary storage and InfiniGuard secondary storage as part of their overall data protection strategy. No performance impact snapshots and highly efficient, low-latency replication provide compelling data protection options. As a purpose-built data protection platform, InfiniGuard delivers high-performance backup ingest, high-speed data recovery, and cyber resilience. InfiniSafe guaranteed cyber recovery is an option for both InfiniBox and InfiniGuard platforms at no additional cost. Finally, IT organizations are assured by Infinidat's guarantees and 100% availability, performance, and cyber recovery.

Infinidat's data protection capabilities are driving down the total cost of data protection, streamlining IT operations, and delivering unprecedented guarantees. Local, replicated, and remote data protection provides the breadth of technical capabilities to meet the varied protection and compliance requirements within an organization as well as removing the complexity in delivery of such a broad set of capabilities.

**InfiniBox provides organizations with the greatest uptime and the highest data availability for the lowest overall costs. With Infinidat, there is no need to compromise.**

Disaster recovery and business continuity are fundamental to IT operations. Infinidat also believes that these fundamental capabilities and features should be a part of the overall storage platform. Considering this, Infinidat provides these technologies at no additional cost. Infinidat is an all-in solution with all software included. Superior capabilities and better TCO are compelling reasons to evolve your data protection infrastructure. Infinidat provides enterprises with a much better overall total cost of ownership when it comes to using the InfiniBox and InfiniGuard as well as help them achieve much greater RPOs, RTOs, VROs and GROs. From an OpEx perspective, having a much smaller footprint to manage, as well as lowering all the environmental (power and cooling) of the data center, have significant advantages. The CapEx is much lower as well, with fewer devices to manage as well as no software licensing costs. InfiniBox provides organizations with the greatest uptime and the highest data availability for the lowest overall costs. With Infinidat, there is no need to compromise.

\* Overall system performance with two DDEs in use, 90% deduplication rate.

\*\* Effective capacities. Actual results may vary.