

WHITE PAPER

Data Protection: A Comprehensive Approach to Protecting Data and Ensuring Availability

Abstract

As data has taken on the role of being the lifeblood of business, protecting that data is the most important task IT has in the data center today. Data protection has many meanings. It ranges from the data availability on primary storage, to the long-term preservation and retention of data for multiple business purposes. In addition, the value of data is getting more and more important each day and there is more of it being generated every second. Protecting that data at greater intervals is becoming more important. This is why the art of protecting data has evolved significantly over the past 30 years.

The requirements, however, for protecting data still have the same baseline. Clients want their data when they want it and where they need it. In order to deliver on these requirements as data scales, some things need to change. This paper will talk about how Infinidat's innovative storage software architecture provides seven nines (99.99999%) data availability from the moment data lands in the array; its unlimited, space-efficient, incredibly fast snapshot capability for quick operational recoveries; its replication engine that provides less than a four-second RPO, and its integration with VMware that provides rapid application recovery. Additionally, this paper will discuss how organizations have replaced multiple purpose-built backup appliances and used InfiniBox® as a backup target due to its incredibly high reliability, scale and high performance for half the cost of traditional backup architectures.

Introduction

The growth of data continues to make it hard to keep up with meeting backup SLAs. Backup windows keep shrinking to the point where we ask ourselves now, “How do you backup a petabyte of data?” The answer, “You don’t!” (Well, the answer, as with most things is never really that simple.) If you tried to protect a petabyte of data by taking advantage of the traditional methods used in the past, such as backup software to disk with a common retention of daily incrementals and weekly fulls, it would be extremely difficult to architect, it would cost an exorbitant amount of money, and you still may never meet your backup windows. Never mind the fact that all of this comes at a time when IT wants to spend less time managing data.

Today, we are finding that all data needs the same level of high protection, or more importantly, availability. You never know what data you’re going to need, when you’re going to need it, or what you may need it for. Add to this, the need to manage both block and file backups, as well as additional demands from the cloud and a whole new paradigm for managing the lifecycle of data needs to be created. Data doesn’t get “stale” anymore, so having the same level of RPO and RTO for all data is becoming more important. Additionally, this is putting a much higher level of stress on IT to be able to meet these SLAs with changing workloads.

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.

Finally, one last thing to consider is that the dynamics of the data center have become increasingly complicated. There are several applications, databases, and file systems that generate different types of data. In addition, as new applications come on line, there need to be ways to protect the newly created data without the need to create a customized backup solution each time. This has led to the creation of multiple data protection silos with multiple products within the data center leaving inconsistent data protection practices and exorbitant costs. IT is currently struggling while trying to get their arms wrapped around these challenges.

Protection Starts with High Reliability

SEVEN NINES OF AVAILABILITY

At Infinidat, data protection starts when the data lands in the array. InfiniBox is designed with seven nines (99.99999%) availability—the equivalent of less than three seconds of downtime per year—a level of reliability no other storage vendor comes close to today. Many competitive solutions can only achieve five nines (99.999%) of availability. The difference between InfiniBox’s seven nines and other storage arrays providing five nines, effectively makes InfiniBox 100 times more reliable, making your data that much more accessible.

INFINIRAID®

InfiniBox protects customer information using a dual-parity, declustered “RAID”-like schema called InfiniRaid®. It 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 method is also known as dispersed data layout and optimizes such data distribution in a 14+2 dual-parity RAID 6-like protection schema. InfiniRaid enables more of every drive in the system to be utilized. This in turn means more usable capacity per system which also helps drive 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.

InfiniBox takes a different approach. Contrary to the clustered protection schema, InfiniRaid rebuild spans the required writes across the entire system; there is no hot spot 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.

InfiniBox’s declustered protection schema 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, in order to ensure 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 schema of the InfiniBox, Infinidat also created an extended metadata checksum that is applied and added to each data block before it is destaged (written) to disk and re-evaluated every time it is read. This eliminates the potential for “Silent-Data Corruption,” which is now commonly seen on larger and denser disk drives, as well as all-flash storage arrays. In addition, each time data is accessed we are constantly looking for this 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.

Snapshots and Replication

IN THE FUTURE, ALL OPERATIONAL RECOVERIES WILL COME FROM SNAPSHOTS


Using snapshots for operational recovery is becoming more and more commonplace. This doesn’t mean that data isn’t “backed up” to tape or disk for longer term retention. However, in order to meet the necessary backup and recovery windows, snapshots are becoming more mainstream.

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 system, with its patented innovative snapshot technology, is designed to provide higher availability, not only for active data, but also for snapshot data. In addition, InfiniBox’s snapshot and replication capabilities help provide RPOs, RTOs, VROs and GROs that are unmatched in the industry for uptime and data availability. VROs are Version Recovery Objectives, and GROs are Geographical Recovery Objectives, which are perfect for snapshots (multiple snaps) and replication (geography). Infinidat snapshots provide IT with the ultimate point in time.

Today’s storage systems spend a non-trivial amount of time “locked.” In other words, these systems are 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; a snapshot is the creation of pointers and inserting them into our trie tree. 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.

Snapshots/Replicas for All Operational Recovery



In order to guarantee the integrity of the snapshot without doing any locking, in the 4K of metadata that is associated with every data section, we include a timestamp. We have timestamps for every 64KB of data that comes into the Infinidat system. Because of this, any write is unambiguously in or out of a snapshot based on the timestamp of when the snap was created and each snapshot has a timestamp. This allows us to have data integrity across all of our snaps without locking any metadata.

Today, we support 100,000 snapshots with zero impact to performance. The 100,000 number has been tested in the lab, however, there is no real theoretical limit. If a client should need more, it's simply a matter of testing. In addition, a volume with or without snapshots has indistinguishable performance. The system can perform as many as 25 snapshots per second and maintain razor flat latency. If clients were to try this on any 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, take snaps of EVERYTHING and perform all operational recoveries from a snapshot delivering much higher risk mitigation. In addition, mounting a snapshot or replica and running off a clone can happen 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. We have clients who take hundreds or thousands of snapshots that take no space and have no performance impact.

There is also additional value with snapshots in InfiniBox. InfiniBox has read/write snapshots on the array that can be used to provide faster and greater access to common data sets for multiple lines of business. Businesses can use these snaps for things such as DevOps, analytics, or reporting, which are giving businesses new ways to have better access to their data. It is also enabling businesses to use less overall storage capacity and provide better data management and greater uptime. Combine this with replication, and lines of business in other locations can have access to the latest and greatest data sets as well—simply and easily.

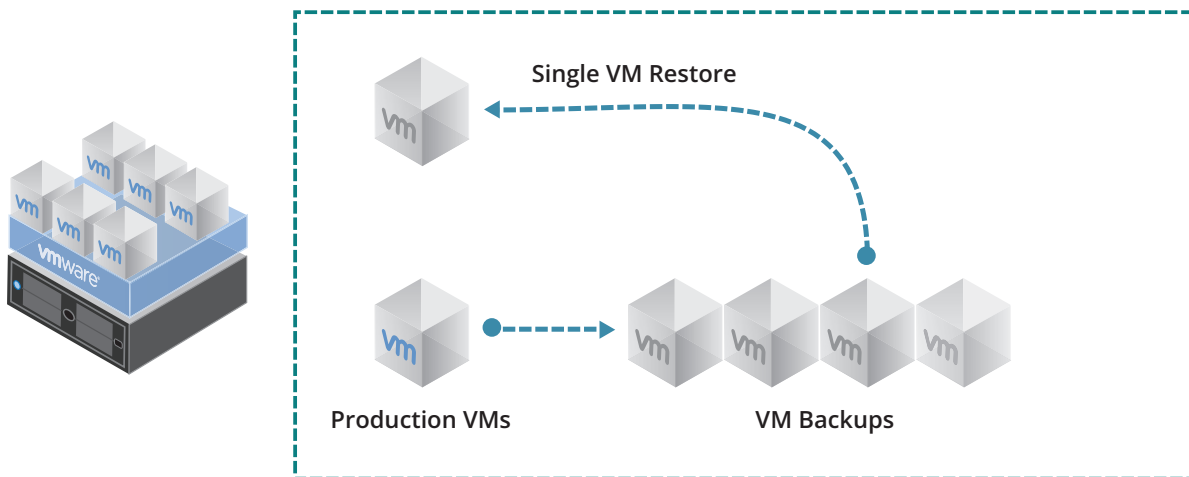
HIGHLY EFFICIENT REPLICATION

Infinidat built its replication engine on top of an 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, the InfiniBox takes a snapshot every n-number of seconds and ships and applies that log at 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.

VMware Data Protection

GREATER AVAILABILITY FOR VMWARE

Another challenge organizations tend to have with their VMware data protection is the growing silos of data protection environments, as well as the growing costs of data protection. VMware has allowed businesses the flexibility to deploy servers for any business purpose quickly and easily, but as these servers move from the testing phase into the production phase they need to be protected. Some groups choose to buy more licenses of the software that is being used and because this tends to be licensed by virtual servers, the costs can become overwhelming as there can be hundreds of servers created. Sometimes when organizations see the expense, they look to other vendors who provide cheaper, and maybe more focused solutions around virtual server backup. This may lower the cost, but it creates silos of data protection practices that may not be in-line with the corporate standards. These silos of data protection environments end up costing the business more money in the long run due to difficult management and data loss.



InfiniBox has done extensive work to integrate with VMware's data protection APIs, such as the vStorage API for Array Integration (VAAI) and the vCenter plug-in, to enable a simple, self-service, single-instance backup and restore solution (without any dependency on vVol or third-party tools). We have also integrated with SRM which takes advantage of the InfiniBox replication and VMware management to enable fast site recovery management. The entire data protection process can be fully automated via the InfiniBox snapshot tool, which maintains consistent copies of thousands of VMs with small retention schedules. InfiniBox snap operations are done via VAAI, utilizing only storage resources. The actual snapshot taken per VM using Infinidat's snapshot technology has zero impact on performance and allows for an unlimited number of snapshots. The Infinidat VMware integration reduces the days and hours for traditional backups and restores of single VM instances down to minutes or even seconds, while also giving organizations better RPOs and RTOs both local and remote. Clients struggling to backup and restore their always growing VM environments can immediately benefit from the powerful InfiniBox snapshot technology. VM level restores are very simple and can be even delegated to the VM admin.

By using snapshots and replication built into the InfiniBox array, clients can exceed all of their data protection SLAs—without additional costs. No longer do you need to buy expensive software licenses or find other solutions to do your operational recoveries for VMware, when you take advantage of the InfiniBox snapshot and cloning capabilities.

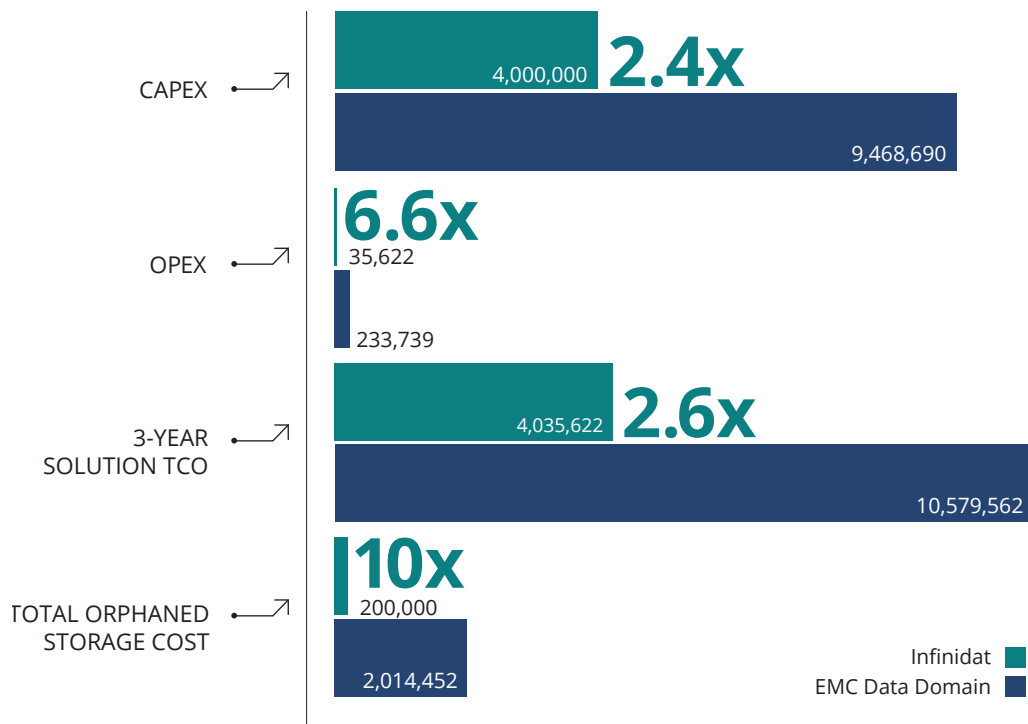
Savings and Efficiencies

INFINIBOX, NOT JUST FOR PRIMARY STORAGE

One other challenge that clients run into as data scales, is the purpose-built backup appliance or PBBA sprawl. As data grows, so does the backup environment. A typical backup environment is four times the size of the primary storage environment. Because of this and the need to keep up with backup and recovery SLAs, clients have purchased a number of PBBA systems (disk arrays for backup) and they are now taking over the data center. The power, cooling and floor space as well as the people required for these systems is a majority of the IT budget. The InfiniBox provides significant capacity per rack with a great deal of efficiency, high performance and low environmental. These allow InfiniBox to drive a low TCO as a data protection target. Organizations who are doing incremental forever backups have seen that our data compression (implemented in version 3.0) provides almost the same value as data deduplication and compression, but at a fraction of the costs. In addition, these organizations that take advantage of the deduplication provided by the backup software are seeing tremendous savings as well.

InfiniBox is used as a backup target by a number of organizations. Because of the scale, reliability and performance of the system, one client replaced 22 EMC Data Domain PBBA systems with just two InfiniBox systems. That is significantly less floor space, power and cooling needed for the data protection environment. In addition, it is fewer devices to manage and fewer things to break and go wrong. The availability of their data protection environment increased by a factor of two by eliminating so many systems. In addition, not only did their OpEx costs go way down, so did their CapEx as they had fewer systems and less maintenance.

Savings from consolidating 22 EMC Data Domain PBBA Systems down to two InfiniBox storage systems.



While today, the InfiniBox does not have all of the storage efficiency features that some PBBAs have, such as data deduplication, the ability to put multiple petabytes of effective capacity in a single 19-inch rack adds a much greater level of efficiency. Organizations can also leverage the deduplication capabilities of their backup software—a much more efficient way to protect data and push it over the network.

Conclusion

Organizations have very complicated data protection environments which cannot be fully covered or satisfied by a single backup schema. Backup today is a mix of many things—different types of disk (PBBAs), different types of tape and different backup software. The flexibility and the shared resource model offered by InfiniBox is very meaningful for cost savings and efficient business operations. It's important to remember that the requirements for backup are only getting more and more strict. Organizations need to be able to access their data 100% of the time, and when they can't, their expectations on how long it should take to recover that data are growing even more demanding.

Today's arrays provide data protection in the form of RAID. However, these RAID capabilities only provide five nines of availability. Next generation InfiniRaid provides even greater uptime in the form of seven nines of availability. In addition, many of today's storage arrays do provide snapshot functionality, but with a capacity and performance impact that renders them not very useful. Snapshots and replication are changing the data protection paradigm as data continues to grow. Next generation snapshots and replication technologies like those in InfiniBox allow companies to protect more data than possible with traditional methods. Performing operational recoveries from snapshots is becoming more commonplace.

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.

Built-in replication also helps businesses to reduce the risk of local disasters and enables businesses to have the right data in the right location, as the business requires. Disaster recovery and business continuity can now be a part of every storage purchase.

Infinidat also believes that all of these features and capabilities should be a part of the overall storage platform. The value of today's next generation storage systems is in the software and its ability to provide the highest available uptime with seven nines of availability, snapshotting and replication all built into the system for one low cost.

Finally, organizations have used the InfiniBox as a target device for their data protection software due to its ease of use, large capacity in a single rack, high availability and high performance for data recovery. Some organizations have consolidated as many as 22 purpose-built backup appliances into two InfiniBoxes.

Data protection is an insurance policy, and like any insurance policy you always want the best rates. Better TCO is the best reason 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 helping them achieve much greater RPOs and RTOs, as well as VROs and GROs. From an OpEx perspective, having a much smaller footprint to manage, as well as lowering all of the environmental of the data center, have significant advantages. The CapEx is much lower as well, with fewer devices to manage as well as lower 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.