



*Maximize the Promise of Hyper-Convergence*

# Table of Contents

Introduction.....	3
Storage Challenges within the Software-Defined Data Center.....	3
What is Hyper-Convergence? .....	3
Maxta at a Glance.....	4
Technology .....	4
Solutions.....	4
Hypervisor Agnostic Architecture.....	4
Architecture.....	5
Overview .....	5
Clustering and Unified Namespace .....	5
Maxta Distributed File System.....	5
Data Integrity.....	6
Benefits .....	6
Increase Flexibility .....	6
Simplify IT Management .....	6
Application Defined .....	6
Maximize Savings .....	7
Software Agility .....	7
Use Cases .....	7
Business Critical Applications .....	7
Remote Office and Branch Office .....	7
Test and Development.....	8
Virtual Desktop Infrastructure.....	8
Disaster Recovery .....	8
Product Features and Functionality .....	8
Simplicity .....	8
Maxta Graphical User Interface – MxInsight .....	8
Historical Data .....	9
Maxta Command Line Interface and REST APIs.....	9
Data Protection .....	9
Snapshots and Clones.....	9
High Availability.....	10
Data Layout.....	10
Fault Management and Serviceability .....	10
Policy Management.....	10
Metro Storage Cluster .....	10
Linear Scalability.....	11
Storage Efficiency .....	12
MxCloudConnect.....	12
Upgrade and Patching .....	12
Performance.....	12
Conclusion .....	13

## Introduction

The Maxta® solution is a hypervisor-agnostic, highly resilient storage platform for the virtual data center. Maxta provides its storage solution as the software-only Maxta Storage Platform (MxSP™) and as the validated Maxta MaxDeploy™ Appliances. Maxta's software-centric, hyper-converged solution is transforming the enterprise storage market. It fully integrates with server virtualization at all levels, from the user interface to data management, and supports all possible deployments of virtual data centers, including private, public, and hybrid clouds. Maxta's software turns standard servers into a converged compute and storage solution, leveraging server-side flash and spinning disks to optimize performance and capacity. Maxta's distributed architecture enables shared storage with enterprise-class data services and full scale-out capability without performance degradation. This results in significant cost savings, as well as dramatically simplifying IT.

This paper will cover the evolution of the datacenter and discuss the technology Maxta has developed to address the challenges faced by the modern IT administrator.

## Storage Challenges within the Software-Defined Data Center

The Software-Defined Data Center (SDDC) is a unified data center platform that delivers converged compute, storage, and networking. SDDC leverages commodity hardware and intelligent software to pool resources and manage them in a global manner for improved utilization and efficiency.

While server virtualization has created an efficient infrastructure by enabling aggregation of compute resources (CPU, memory), storage technologies have only made incremental innovations in the past decade. Shared storage, generally SAN or NAS, is a requirement for utilizing many of the benefits of server virtualization. However, these networked storage solutions are expensive and increase the complexity of storage management due to the mismatch between storage constructs (e.g., LUNs, volumes, file systems) and virtualization constructs (Virtual Machines or VMs). The advent of flash storage has significantly improved the performance of storage, but these solutions do not address the issues of cost and complexity.

In recent years, alternative approaches to networked storage, such as virtual storage appliance (VSA) and distributed file system, have emerged. However, these approaches are limited in terms of scalability and performance, while not improving storage management. Thus, they are not an alternative to networked storage in most cases. Other alternatives to networked storage such as Hadoop have emerged as well for Big Data environments and specific storage workloads. However, these alternatives are not well suited for general purpose virtualized environments and associated storage workloads.

## What is Hyper-Convergence?

Hyper-convergence and Software-Defined Storage (SDS) address the challenges of storage in virtualized environments by implementing storage functionality entirely in software while leveraging commodity components. The software is deployed on the servers that are running the hypervisor and enables pooling of disperse server-side physical storage resources. SDS provides VM-centric management of the storage pool, thereby addressing the gap between storage constructs and virtualization constructs.

In this way, a well-defined hyper-converged infrastructure will deliver resource efficiency by utilizing industry standard servers, SSD, PCI-attached flash, and magnetic disk drives. It will provide operational efficiency by managing integrated software entities (converged compute and storage) rather than isolated hardware resources (storage arrays). A software-defined approach also allows for quick responses to support new requirements. Finally, there should be no need to coordinate or negotiate storage resources with a separate storage team and no need to forecast and manage different physical storage resource pools.

## Maxta at a Glance

### Technology

Maxta has developed a VM Storage Platform that enables IT to fully realize the vision of the virtual data center. MxSP addresses the issues discussed above:

- Hypervisor-agnostic implementation currently supports VMware vSphere, KVM, and OpenStack
- Full integration into the server virtualization user interface
- Scale-up and scale-out compute and storage independently on demand
- VM-level storage abstraction rather than block or file level abstraction
- VM-centric enterprise-class data services
- Support live migration, dynamic load balancing, high availability, data protection, and disaster recovery
- Optimized for flash performance and hard disk capacity

In summary, MxSP significantly simplifies IT, increases IT efficiency, and dramatically reduces capital and operational expenditures.

### Solutions

Maxta offers two different types of solutions:

- MaxDeploy Hyper-converged Appliance
- MxSP software-only

The MaxDeploy appliance delivers a flexible way of deploying hyper-converged solutions via predefined and pre-validated solutions that combine Maxta's software along with partner solutions and platforms. This removes interoperability and performance guesswork and simplifies the ordering process.

MxSP software-defined storage solutions provide customers the complete flexibility to customize their solutions and run on existing x86 servers. Maxta's server configuration tools and appliance briefs can help with configuration suggestions and compatibility verifications.

### Hypervisor Agnostic Architecture

MxSP is a hypervisor agnostic implementation of an enterprise storage solution for virtualized environments. The diagram below describes the architectural implementation.

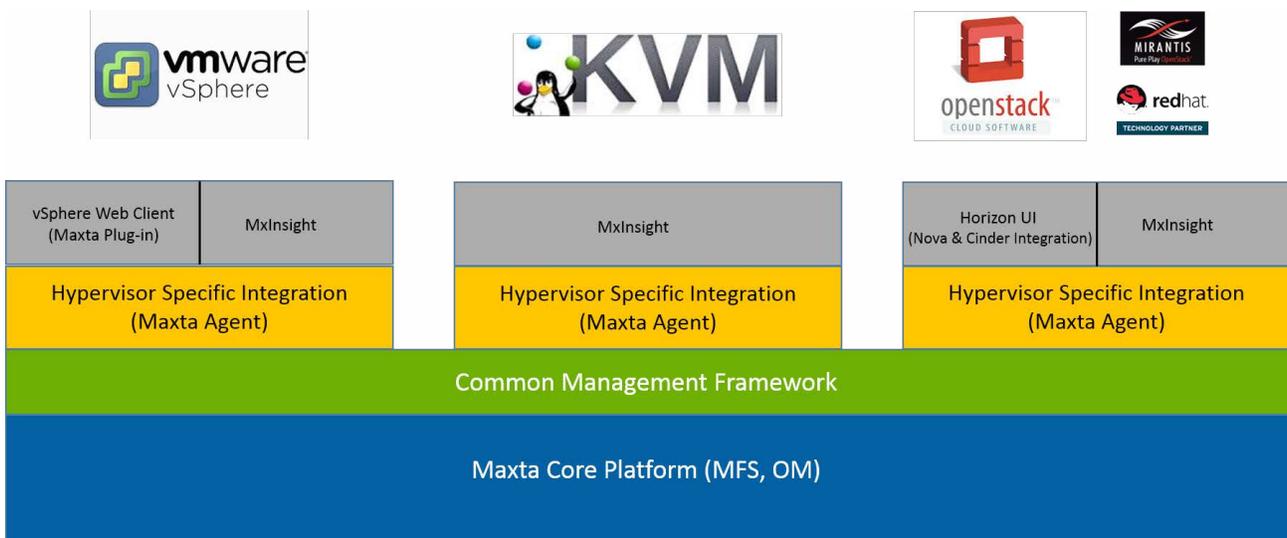


Figure 1: MxSP Hypervisor Agnostic Architecture

The Maxta Core Platform consists of the Maxta Distributed File System (MFS) and Object Manager (OM). The core platform and Maxta's Common Management Framework are common across all hypervisors. A Maxta agent is developed to integrate into each specific hypervisor. The same MxInsight web management interface is available on all hypervisors. Additionally, Maxta has developed a plug-in into the VMware vSphere web client for MxInsight as well as seamless integration into OpenStack's Horizon UI.

# Architecture

## Overview

MxSP's innovative, peer-to-peer architecture aggregates storage resources from multiple servers, incorporating a global namespace and creating a Maxta storage pool. An instance of MxSP software is installed on each of the servers that are part of the virtualization cluster. The software leverages flash storage resources such as PCIe, NVMe, or SATA SSDs for performance and SATA or SAS HDDs for capacity. All the servers running MxSP software have access to the aggregated Maxta storage pool. However, it is not a requirement that all servers have storage resources or contribute storage resources to the Maxta storage pool. It is also not a requirement that the servers that contribute storage resources to the Maxta storage pool contribute the same capacity. The communication between the MxSP instances is over a private Ethernet network. It is highly recommended that the network be a private network of at least 1GbE. The servers running MxSP either contribute storage to Maxta storage pool (converged compute/storage servers) or do not contribute storage to the Maxta storage pool (compute-only servers).

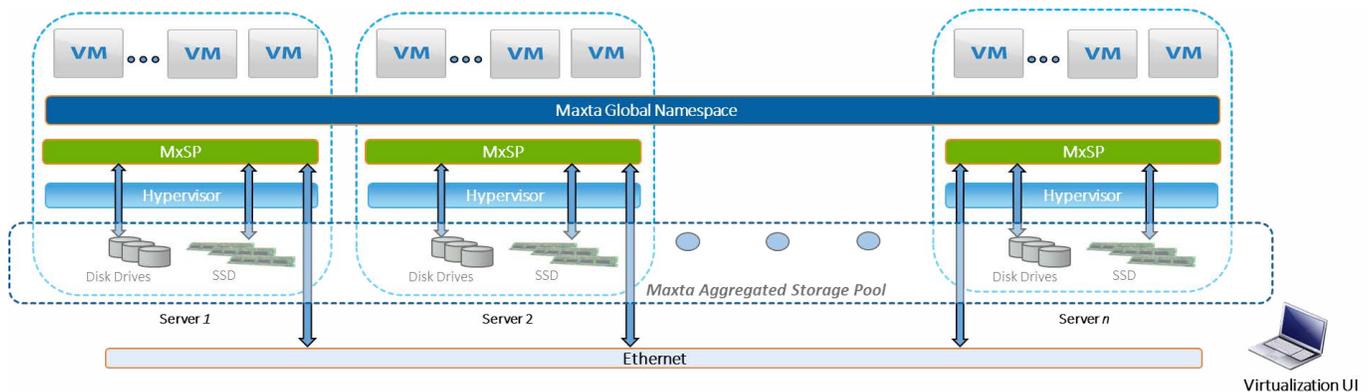


Figure 2: MxSP Architecture

There are a few aspects of the Maxta architecture which have been specifically optimized for each hypervisor. In VMware environments, Maxta storage is presented as an NFS datastore shared across all ESXi hosts within a cluster. MxInsight is directly integrated into the vSphere Web Client to provide full management capability from a single pane of glass.

Maxta has developed drivers for Cinder (block storage) and Nova (compute) to support OpenStack deployments. These plug-ins enable administrators to create thin-provisioned instances and volumes as well as time, performance, and capacity-efficient snapshots and clones. These operations can be performed directly through the Horizon UI. Additionally, the same MxInsight interface that exists for VMware deployments is available for KVM and OpenStack environments. MxSP can also be run in a KVM environment without OpenStack.

## Clustering and Unified Namespace

MxSP presents shared storage at cluster-level visibility. The unified namespace service provides all nodes in the Maxta cluster with visibility to the same hierarchy of objects. This service also provides all nodes with the ability to modify any of these objects. At any point in time, a minimum of three nodes must be configured to run the unified namespace service, and a majority of those nodes must be online to maintain quorum. MxSP's high availability services ensure that data remains accessible when nodes fail and can provide greater fault tolerance depending on cluster size, replica policy, and fault containment units made available via Rack Awareness.

## Maxta Distributed File System

MxSP uses Maxta's own highly scalable, resilient Maxta Distributed File System (MFS) to deliver shared storage with enterprise class services. MFS has been developed with the purpose of meeting the needs of an enterprise hyper-converged virtual environment. MFS eliminates split-brain conditions by synchronizing operations between instances across the cluster. MFS runs on each server in the virtualization cluster to deliver a highly available and scalable global namespace.

## Data Integrity

MxSP uses strong checksum algorithms to prevent against silent data corruption in a virtualized infrastructure. Maxta protects against data corruption scenarios by performing checksums at multiple points in the data path. Data integrity is checked whenever the data is stored to a disk, whether that be on flash or when data has been de-staged to spinning disk, and anytime the data is sent across the network. Checksums are validated each time data is read. If an error is found, MxSP repairs and recovers the data.

## Benefits

### Increase Flexibility

Maxta delivers on the promise of hyper-convergence by providing flexibility and choice at all levels of the datacenter. MxSP's hypervisor agnostic architecture allows administrators to pick the hypervisor of their choice. The software-defined approach provides flexibility when it comes to servers and disks, enabling customers to deploy Maxta's hyper-converged solution on the hardware of their choice. Maxta leverages flash media installed on commodity servers and disks to maximize performance while minimizing cost. MxSP eliminates the vendor lock-in commonly associated with appliance based hyper-converged solutions. Additionally, Maxta provides flexibility in purchasing the software through a one-time upfront perpetual license or a recurring subscription-based service.

### Simplify IT Management

MxSP dramatically simplifies storage management by eliminating the need for provisioning storage or managing volumes, LUNs, file systems, RAID, ACLs, etc. This removes the necessity to develop the mapping between virtual machines and storage constructs such as LUNs and volumes, minimizing administrative errors. Converging compute and storage on standard servers eliminates storage networking tasks such as zoning. MxSP provides a single-pane-of-glass for both VM and storage administration by incorporating all storage functionality within the virtualization UI. Once MxSP is installed, the administrator can simply point to the Maxta storage pool from the virtualization UI during VM creation and MxSP will take all steps to optimally provision and allocate resources for the new VM.

The installation and configuration of MxSP takes only a few minutes, enabling all system/VM administrators to install and manage storage. During installation, MxSP aggregates storage across all the servers in the cluster and presents it as the Maxta storage pool with a global namespace that is used to provision storage for VMs. MxSP will not add any previously allocated storage devices to the Maxta storage pool. It is also possible to selectively exclude any storage device from the Maxta storage pool. The installer automatically installs and configures all the necessary pre-requisite software and identifies the minimum hardware and software requirements for installing MxSP. The installer guides the user through the resolution process in case the minimum requirements are not met.

The Maxta storage pool and all data services such as replication, snapshots, and zero-copy clones are configured and managed from the VMware vSphere Client at VM granularity. This simplification eliminates the day-to-day tasks of storage management and enables administrators to focus on managing applications and VMs.

Administrators can interact with the Maxta Storage Platform through MxInsight (Maxta Graphical User Interface), Maxta Command Line Interface (CLI) and REST APIs.

### Application Defined

Traditional storage arrays offer the ability to define storage on a per-application basis through LUN configurations. However, this is a very complex process which requires investment in heavily specialized storage administrators. Hyper-convergence has relieved this pain by simplifying management and deployment. Unfortunately, by creating a single, uniform pool, other hyper-converged solutions have lost the flexibility to configure storage for each application – one size is NOT good for all.

Maxta brings the best of both worlds with its application defined storage platform providing both ease of management and the ability to customize storage to individual applications. Through Maxta's management interface, users can easily set policies with just a few clicks on a per-VM or per-vDisk basis. These policies help to configure each application for failure tolerance and data layout, among other characteristics.

## Maximize Savings

MxSP delivers the capabilities of storage arrays in software by leveraging server-side SSD and internal drives or JBODs connected to standard servers. MxSP enables significant CAPEX savings by converging compute and storage resources on standard servers leveraging commodity components, without compromising performance or scalability.

MxSP eliminates all storage provisioning activities and dramatically simplifies day-to-day data management activities such as recovery of virtual machines from snapshots and creation of new VMs rapidly using clones. Additionally, MxSP eliminates storage management tasks such as changing RAID parameters or storage tier management. MxSP provides significant reduction in power, cooling, and floor space by converging compute and storage resources on standard commodity servers and eliminating the need for storage arrays.

### Software Agility

The traditional storage model packages hardware together with non-transferable software. When storage hardware is refreshed every four to five years, software must be purchased again. Hyper-converged appliances are no different in this model of non-transferable software. In fact, this approach forces end-users to make a compromise due to the faster pace of innovation in compute technology compared to that of storage. Adopters of hyper-converged appliances must either: refresh compute at a slower pace (resulting in hardware which is out of date and slower renewal rates for partners) or refresh storage at a faster pace (creating a higher storage acquisition cost for the customer).

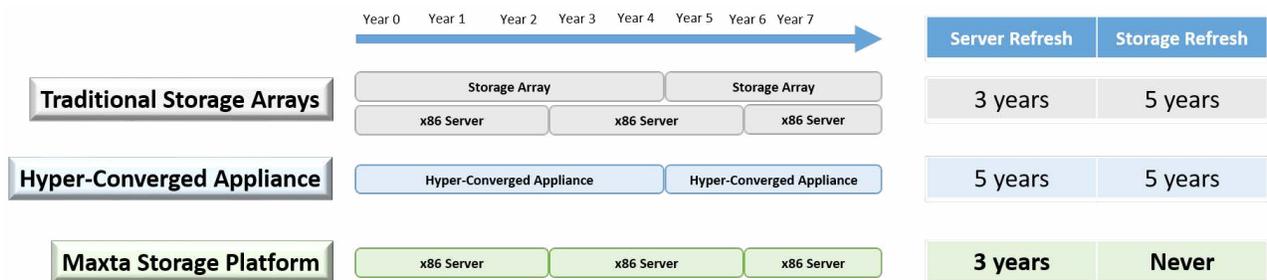


Figure 3: Software Agility

Maxta de-couples the software from the hardware with its fully transferable software license. This enables end-users to purchase the hardware they need, when they need it, without having to repurchase storage software – Maxta is the last storage software purchase in the datacenter.

## Use Cases

The Maxta solution can be provided directly as a solution to enterprises or Managed Service Providers (MSPs) looking to increase their margins. MxSP integrates into existing frameworks for user access control and multi-tenancy, thereby allowing MSPs to enforce resource constraints as desired throughout their environment.

## Business Critical Applications

The flexibility of the Maxta Storage Platform enables it to support a variety of use cases in the data center. MxSP can be the storage solution for business critical applications. MxSP's full suite of enterprise class services meet the criteria necessary to support workloads running on primary storage.

## Remote Office and Branch Office

MxSP is also a great fit for Remote Offices and Branch Offices (ROBO). ROBO environments require the same enterprise-class data services of primary storage, but operate within a more limited budget and face more infrastructure constraints. Maxta's focus on VM-centric management overcome the lack of specialized skills commonly found in ROBOs. Additionally, bringing compute and storage together onto hyper-converged nodes greatly reduces space and power requirements.

## Test and Development

Maxta's zero-copy snapshots and clones are advantageous to test and development teams who need point-in-time copies of the latest production data. Rather than copying an entire LUN, as is the model for traditional storage arrays, Maxta's snapshots and clones are taken at the VM-level. The snapshots and clones are also time-, performance-, and capacity-efficient, meaning that they can be created and deleted instantly, have no impact on performance, and do not take up any space on creation.

## Virtual Desktop Infrastructure

Maxta's snapshots and clones are also highly beneficial to Virtual Desktop Infrastructure (VDI) deployments in which administrators have to deploy thousands of virtual desktops in minutes. These environments have to maintain a better end-user experience than a traditional stand-alone desktop at lower cost. MxSP is able to meet the performance requirements of VDI by leveraging SSDs as read and write cache while providing a cost-efficient solution.

## Disaster Recovery

Maxta storage can serve as an endpoint in Disaster Recovery (DR) environments. MxSP leverages the existing replication software (Zerto, Veeam, etc.) to provide VM-level recoverability. The recovery site maintains all of the enterprise-class capabilities of MxSP, including the ability to maintain short term backups via Maxta's snapshots.

## Product Features and Functionality

### Simplicity

#### Maxta Graphical User Interface – MxInsight

Maxta's management console, MxInsight, is a responsive web application which can be accessed in two ways: a single pane of glass integrated into the hypervisor's virtualization user interface or its own webpage. MxInsight provides visibility at VM-level granularity, with no need to manage LUNs or other storage constructs.

The management console has been designed to provide administrators with more visibility into the status and performance of the Maxta cluster. The dashboard displays historical and real-time statistics related to capacity, latency, IOPS, and throughput. There are also statistics for each node, virtual machine, and disk to give administrators further insight into the performance of their cluster. MxInsight provides disk information on a per-node basis, allowing administrators to view a breakdown of storage pool capacity usage for each disk in the Maxta cluster, including SSDs and flash media.

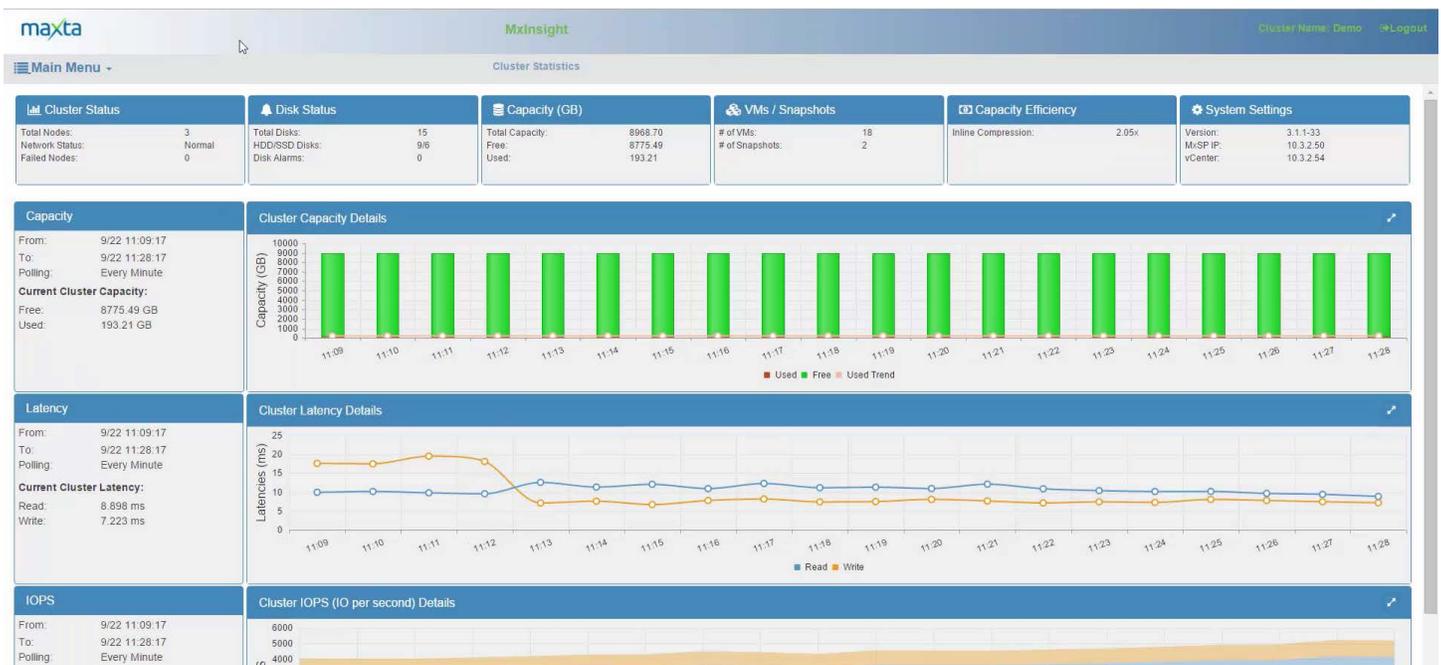


Figure 4: MxInsight Dashboard View

Administrators can monitor events, alarms, and recent tasks through MxInsight, as well as view the status of ongoing tasks. MxInsight can be configured to trigger alerts at different thresholds for each indicator, including SSD wear. There are also options to configure language preferences and email notifications via SMTP.

MxInsight streamlines the workflow of service actions such as adding and removing disks and hosts from the Maxta cluster. Maxta's user interface allows administrators to create and manage VMs with the standard set of features available in the hypervisor manager, as well as policy management functionality unique to MxSP. Individual VM and virtual disk policies can be configured during VM creation and modified after VM creation, whether that VM was created through MxInsight or the hypervisor management console. Administrators can set policies for creation and retention of snapshots at specified time intervals, and they can automate the creation of multiple clones with customization specifications.

## Historical Data

MxSP uses a statistics database to maintain historical data points for the following metrics as they relate to the Maxta storage cluster:

- Used and free capacity on cluster, node, and storage pool levels
- Read and write latency on cluster, node, storage pool, virtual machine, and virtual disk levels
- Read and write IOPS on cluster and node levels
- Read and write throughput on cluster and node levels
- Percentage of cache hits at VM and virtual disk level, broken down into MxIL, in-memory, read cache, and misses

## Maxta Command Line Interface and REST APIs

The Maxta Management Server provides an access point into the Maxta CLI and REST APIs. Administrators can log in to the management server as the console user and directly interface with MxSP using the Maxta CLI. Maxta's REST APIs are exposed via the Maxta Management Server for all features of MxSP. For further documentation of the Maxta CLI and REST API, please see the Maxta Technical White Paper and Maxta User Guide.

## Data Protection

### Snapshots and Clones

MxSP delivers an unlimited number of time, performance, and capacity-efficient zero-copy snapshots and clones. These snapshots and clones are configured and managed at the VM level rather than at the storage level. This enables the VM administrator to leverage advanced storage features without the need for deep storage and vendor specific expertise.

Maxta's zero-copy snapshots are read-only, point-in-time virtual copies of data which do not consume any capacity and do not require any upfront space reservation. Snapshots are created in less than one second, independent of the capacity of the source, and snapshot creation does not impact the performance of virtual machines. Snapshots can be created even when a VM's replica data is not available due to an offline node. Maxta's snapshots can also be deleted quickly and out of order. Snapshots on MxSP provide the first layer of data protection without sacrificing performance or capacity. MxSP also enables application consistent snapshots to ensure that applications start from a consistent state when recovering a VM to a snapshot's point-in-time.

Maxta's zero-copy clones use the same methodology as Maxta's snapshots with the difference being that clones are read/write virtual copies of the data. Clones do not consume any capacity unless data is modified, do not require any upfront space reservation, are also created in less than one second, and do not impact the performance of virtual machines. Administrators can take snapshots of clones and build out full snapshot/clone hierarchies with the ability to revert and power up VMs. In addition to creating single clones, MxSP supports automated creation of multiple clones with the ability to add guest customization scripts.

## High Availability

### Data Layout

MxSP optimizes the data layout for the typically random read/write operations of virtual workloads. Under normal operation, reads and writes are served by the server hosting the VM. MxSP provides the ability to maximize the probability that a VM's data is read from the node's local cache, eliminating the network round trip latency associated with distributed storage systems.

MxSP both stripes and mirrors (replicates) data in order to create redundancy and maintain availability during hardware failures. The mirroring policy is configured during VM creation and can be set to a different number of replicas at any time. MxSP also performs proactive capacity rebalance to ensure that no node or disk within the Maxta cluster becomes over-utilized.

### Fault Management and Serviceability

MxSP has been developed with a focus on quickly and easily recovering from failures with zero data loss and zero downtime. MxSP has been designed to avoid having any single point of failure in the cluster and to be able to quickly and automatically recover from any failure. The storage platform is able to withstand failures to spinning disk, flash media, servers, sites, racks, and Maxta services without suffering any downtime or loss of data. MxSP's resiliency covers both transient and permanent failures. Additionally, Maxta is able to leverage the hypervisor's built-in high availability protection, such as using NIC-teaming to protect against network failure.

The copies of data maintained in replicas enable rapid recovery from failures, and additional replicas ensure that the datacenter can withstand multiple failures without seeing any downtime. For example, a VM deployed on the Maxta cluster with three replicas of data will persist through two failures without data loss or downtime.

### Policy Management

Maxta's policy management provides further granularity and control over the hyper-converged environment. In keeping with MxSP's VM-centric principles, administrators can set policies on a per-VM or per-vDisk basis. MxInsight provides the ability to create policies upon VM creation and modify policies after a VM has been created, whether the VM was created through MxInsight or the hypervisor's centralized management interface.

When creating a VM through MxInsight, the following policies can be set on a per-VM basis: striping, number of replicas, node affinity, rebuild priority, and metro clustering. These policies can also be modified after a VM has been created.

During VM creation, the following policies can be set through MxInsight on a per virtual disk basis: page size, compression, and read cache. All of these policies can be modified after a VM has been created. When adding a virtual disk to an existing VM, all of the vDisk policies can also be configured.

MxSP's snapshot policy management also enables automated creation and deletion of snapshots, along with retention schedules. Snapshot policies can be configured with a start date and optional end date, a frequency of daily, weekly, or monthly, and a time interval during which snapshots will be taken on a customizable periodic basis. Users can also define how long to retain the snapshots, based either on time period or number of snapshots.

### Metro Storage Cluster

Maxta's Metro Storage Cluster capabilities can be enabled if Rack Awareness is configured during MxSP installation. A metro storage cluster (or stretched storage cluster) delivers workload mobility, cross-site automated load balancing, and continuous data availability across datacenters.

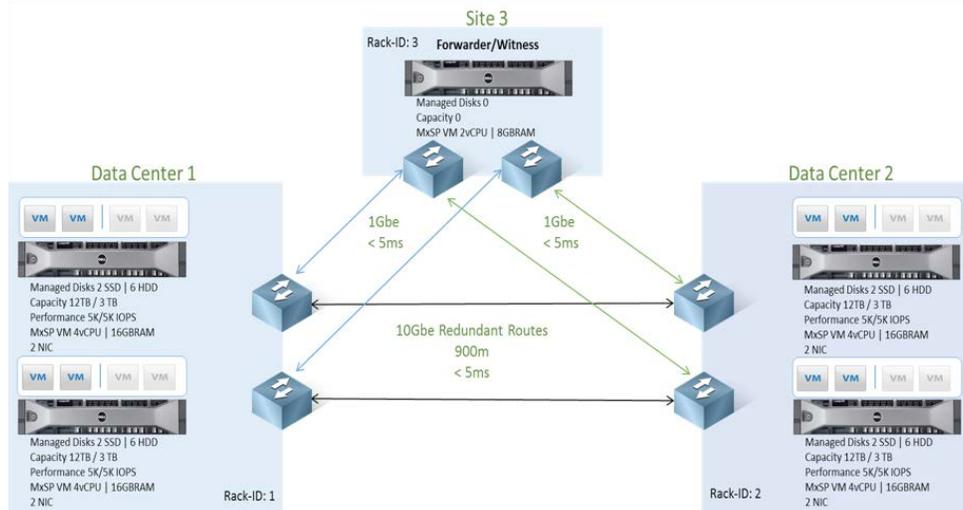


Figure 5: Maxta Metro Storage Cluster

The technology can be adopted to provide continuous availability between racks within a datacenter, across buildings within a campus, or across datacenters spanning different cities. The key requirement to support Maxta Metro Storage Cluster is a maximum latency of 5 milliseconds round-trip time (RTT) across all datacenters for the storage network.

MxSP uses synchronous replication across sites to provide metro storage cluster capabilities. The placement of replica data is performed using the Rack ID's which are configured via the Rack Awareness installation properties. MxSP intelligently performs reads from the node with lowest latency, therefore read requests will most likely be served from the local rack.

### Linear Scalability

The versatile architecture of MxSP provides the ability to scale capacity and performance independently on-demand without having to over-provision resources. MxSP also delivers linear scalability as new servers are added to the cluster without impacting performance. The example below highlights the linear scalability of the platform in a VDI use case. The average latency of accessing the virtual desktops remained almost the same even after increasing the number of VMs deployed and adding nodes to the cluster. The linear scalability of MxSP from 100 to 456 desktops without impacting performance is shown below:

Maxta Storage Platform (MxSP)			
Number of Desktops	100	300	456
Number of Nodes (ESXi hosts)	2+1*	3	4
Number of stuck or unresponsive desktops (Sessions)	0	0	0
Controller VM configuration (CVM)	4vCPU/8GB Memory	4vCPU/8GB Memory	4vCPU/8GB Memory
Performance Metrics			
Minimum response time (VSIbase)	731 ms	797ms	843 ms
Average response time (VSImax Average)	823 ms	929 ms	995 ms
Maximum response time (VSIMax Threshold)	1731 ms	1798 ms	1844 ms

\* 2 Hyper-Converged Compute/Storage Nodes + 1 Compute-Only node

Figure 6: MxSP Scales Linearly in VDI Deployments

## Storage Efficiency

MxSP delivers capacity optimization capabilities such as thin provisioning, in-line compression, and in-line de-duplication to enable significant storage efficiency. Thin provisioning provides the ability to provision capacity several times larger than the amount of physical storage resources. This enables procuring incremental capacity on demand. The combination of compression and thin provisioning provide the average Maxta deployment with 2-4X in capacity savings.

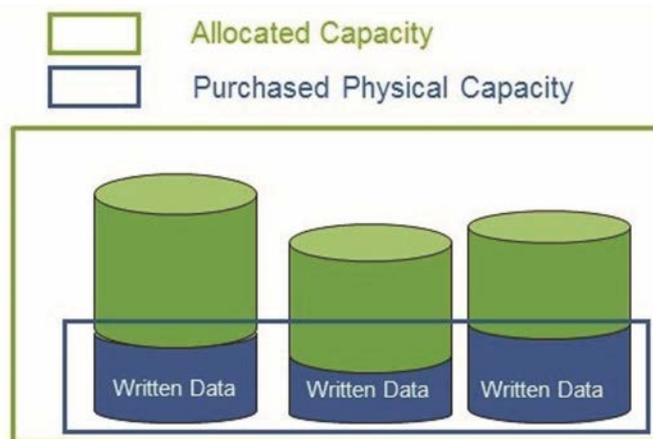


Figure 7: MxSP Maximizes Capacity Savings

## MxCloudConnect

Maxta has developed MxCloudConnect, a cloud based proactive reporting and alerting system which is included as part of the standard MxSP installation. This service provides administrators with access to a web-based portal which has information on alerts and events across multiple Maxta clusters. MxCloudConnect can also be used by managed service providers to monitor client deployments of MxSP and ensure that all clusters are functioning properly. Maxta uses MxCloudConnect to check on the state of all MxSP deployments and respond immediately to any alerts. This quickens response time on support calls and reduces customer headaches.

## Upgrade and Patching

MxSP performs a fully automated rolling upgrade wherein the new bits are applied to each Maxta controller VM in the cluster one node at a time. There is no need to move virtual machines off the servers in the cluster – patch and upgrade of MxSP is completely non-disruptive and requires zero downtime. The process for upgrading the hypervisor is very similar to that of upgrading MxSP.

## Performance

Maxta has optimized its I/O path to maintain high performance in virtualized environments which utilize a hybrid disk configuration consisting of spinning disks and flash media. Read and write operations are directed to the flash media or SSD before spinning disk to maximize performance.

Data is written to the flash media in a sequential manner and replicas are written to flash on other nodes in order to maintain data availability. The writes are later de-staged to magnetic drives in a sequential manner. Performing writes to flash and magnetic drives in a sequential manner maximizes I/O performance and SSD endurance.

Read operations are accelerated by maintaining metadata and frequently accessed data on the SSD cache. MFS also leverages RAM from its controller VM as a cache tier ahead of flash. While MFS has been optimized to deliver cost efficiency and high performance in hybrid storage configurations (mix of HDD and flash/SSD), all-flash or all-SSD configurations are also supported.

A best practice for MxSP deployments is to have two flash media devices (eMLC SSD or PCIe/NVMe flash) per node to support mirroring for failure scenarios. Maxta recommends that 5-8% of HDD capacity be available in flash. These numbers double in environments where SSD mirroring is enabled.

## Conclusion

Maxta is a hypervisor agnostic implementation of software defined storage supporting VMware vSphere, KVM, and OpenStack. The increased flexibility enables customers to select their server of choice. Maxta simplifies administration and minimizes the learning curve by integrating into the virtualization management interface. MxSP dramatically simplifies IT by eliminating the need for provisioning storage and managing volumes, LUNs, file systems, and RAID. The installation and configuration of MxSP takes only a few minutes. Additionally, all data services, such as replication, snapshots, and zero copy clones, are configured and managed at the VM level rather than the storage level. This enables the VM administrator to leverage storage without the need for deep storage and vendor specific expertise. With MxSP, administrators can manage VMs and not storage.

MxSP delivers enterprise-class data services with best-in-class resiliency, continuous availability, data protection, and agility. With unlimited VM level snapshots, backup performance is improved and recovery times are dramatically reduced. By leveraging VM level clones, rapid provisioning of VMs is possible. With continuous data protection, disaster recovery of applications is greatly simplified and accelerated for business continuity. MxSP seamlessly integrates with all the advanced capabilities of the server virtualization software.

Maxta's MaxDeploy appliances deliver a predefined and pre-validated solution that removes interoperability and performance guesswork and simplifies the ordering process.

MxSP enables significant capital savings by converging compute and storage resources on standard commodity servers, without compromising performance or scalability. This provides considerable up-front capital savings compared to storage arrays. In addition, MxSP leverages any combination of magnetic disk drives and flash technology, snapshots, zero-copy clones, thin provisioning, in-line compression, and in-line de-duplication to increase storage efficiency and reduce storage expenses. By significantly simplifying IT, increasing IT efficiency, and enabling administrators to focus on managing applications and VMs, MxSP enables dramatic reduction in operating expenses.



Think Outside the Storage Box

[www.maxta.com](http://www.maxta.com)



[@MaxtaInc](https://twitter.com/MaxtaInc)



[www.youtube.com/c/MaxtaInc](http://www.youtube.com/c/MaxtaInc)



[www.linkedin.com/company/maxta-inc-](http://www.linkedin.com/company/maxta-inc-)



[www.facebook.com/maxtacorp](http://www.facebook.com/maxtacorp)

© 2015 Maxta. All rights reserved. Maxta, MxSP, and MaxDeploy are trademarks of Maxta Inc. All other trademarks are property of their respective owners.

[sales@maxta.com](mailto:sales@maxta.com) | 1-844-4-4MAXTA