

MatrixStore Explained



ObjectMatrix



Concept

MatrixStore is a software product focused on protecting digital assets and making them available on-line and on-demand. MatrixStore adds scalability, security, search and self-management to commodity disk-based storage hardware.



The MatrixStore software is not tied to proprietary hardware; it utilizes high quality, off-the-shelf disk-based storage hardware. MatrixStore server nodes are run, and perform all compliance requirements, on the Mac OS or Linux and therefore can support hardware platforms that are supported by those operating systems.

Clients can connect from all key operating systems.

Product Highlights

- distributed Index and Search
- self healing, self managing clustered architecture
- 5 "9's" data availability using commodity hardware
- open API for application integration
- low management overhead
- highly competitive TCO
- enhanced meta-data and business rules support
- scalable in both capacity and performance
- 256 bit encrypted access
- independently audited for security by e-security specialists Ubizen
- non-proprietary storage format
- access to all data, all of the time

Reference Data

MatrixStore is suitable for all reference or fixed data, data which will not change. This is the type of data that has seen impressive growth over the last couple of years with the majority of analysts predicting the trend to continue a pace.

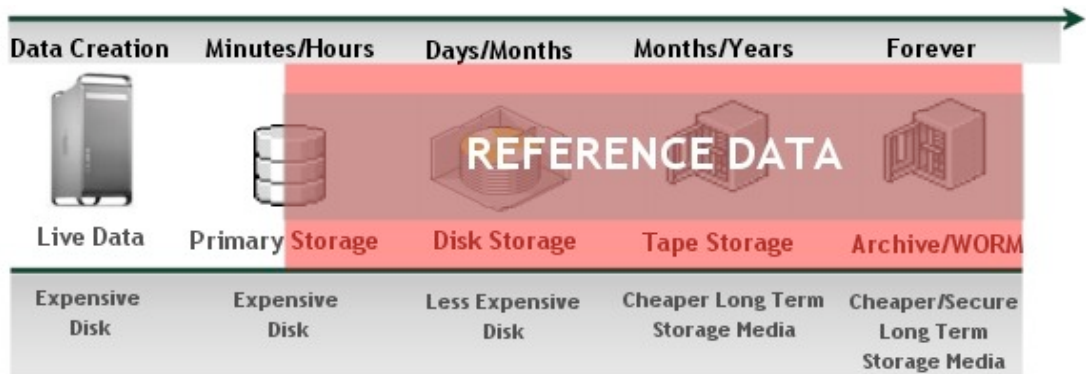
- Movie Files
- Post production rich media
- 3d seismic videos
- Satellite Image Files
- MP3 Music Files
- Medical Records
- Photography data
- Scientific data sets
- MRI data
- Email
- PDF reports



Positioning

Traditional Tiered Storage

For the majority of its lifetime data is fixed, it is highly unlikely to change after it has been created. Fixed or Reference Data is traditionally stored against a tiered architecture that is expensive to maintain and by its nature reduces the availability and thus the value of the data being stored.



Often expensive disk such as a SAN environment is used to keep reference data available despite the inherent costs involved. The drop in disk price also seems to correlate directly with such architectures being used more and more inappropriately. Once the administrators have started to move the data through the tiers that data is taken off-line and is, as such, useless to anyone.

Digital Asset Archiving

MatrixStore is designed to keep reference information alive and available throughout its lifetime. As you can see below the SAN-like environments can be kept to a minimal, more effective and manageable size by migrating data onto MatrixStore the expensive disk is allowed to do what it does best, fast uncluttered access to data.

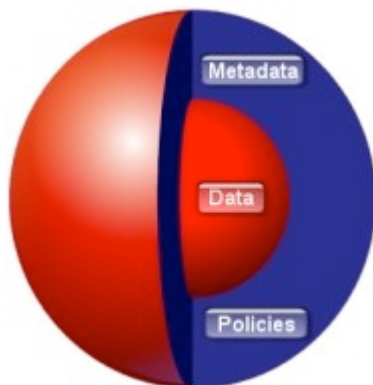


An organisation benefits from reduced operating costs brought about by reducing the management and expertise required to keep the data in a tiered architecture whilst leveraging the value in their digital assets that will now be available 24/7.

Softer, less tangible benefits arise from increased productivity both in terms of quicker, more reliable restores when required and members of the team having the data they require at their finger tips.

MatrixStore Objects

All data archived to MatrixStore is stored as 'Objects'. Each Object has 3 components; the data, meta-data and policies.



Data

This is the actual data itself, as it would normally be stored on any other file system. An important fact to note here is that MatrixStore does NOT change the format of the data on the disk.

MatrixStore does **not** change, alter or store your data in a proprietary format.

Meta-data

This is descriptive information about the Data. It could be the artist, album, title, length or genre information about a music file, or could be co-ordinates, type, range, date etc. about a satellite image file. This meta-data can be any key/value pair that the application deems suitable and useful for object retrieval at a later date.


It is also possible to update attributes on an object at a later date which is useful to add extra context to an asset over its lifetime.

Policies

This is the information about how the data should be stored, such as the number of copies, the retention period and the permissions.

Location Independence

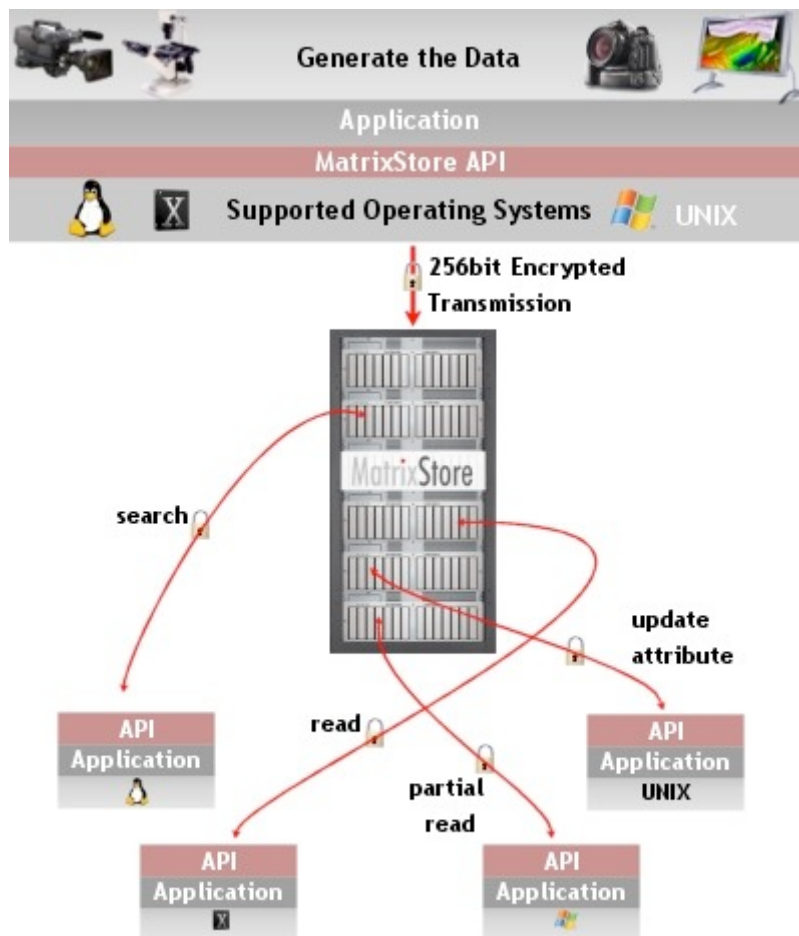
Treating the data as objects allows MatrixStore to hide the management of the location of data within the cluster. MatrixStore automatically finds the most appropriate location in the cluster to store an object. Applications just talk to the MatrixStore like a black box; they send data to it with no knowledge of where the data eventually resides. Upon completion of a successful write operation the application receives a globally unique Identifier with which it can retrieve the data at any time. MatrixStore does NOT address the data by its contents.



Accessing MatrixStore

All access to MatrixStore is via a powerful API through which all data transmission is 256bit encrypted meaning the data can be accessed by authorised personnel wherever there is an Internet connection which reduces the costs of providing VPN access. Users of MatrixStore will access data using an application that has integrated to the MatrixStore API such as Artbox a digital asset management solution from Proximity Group.

MatrixStore is delivered with "DropSpot", an interactive archiving, search and restore desktop application that runs on Windows, Linux and Mac OSX. DropSpot allows direct use through the GUI and also allows scripting to quickly archive data in an existing work-flow process.



Data is first written to MatrixStore using the API. Once data resides in MatrixStore it can be retrieved by any authorised client be it by search or knowledge of the objects unique Id.

The MatrixStore API is available in both C and Java. A typical integration with the API should take less than a week development time plus any testing cycles. Please refer to the MatrixStore API User and Reference Guides for more detailed information.

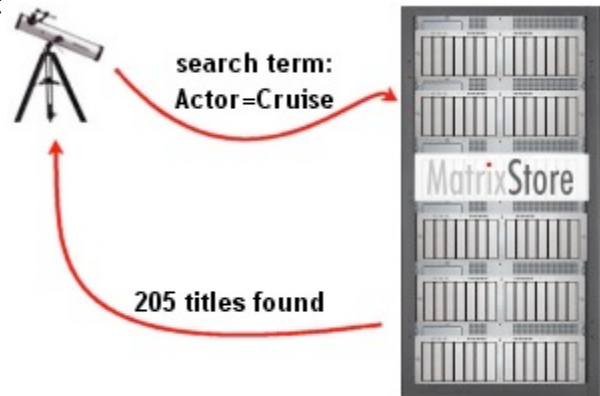
Functionality

Search

MatrixStore provides full, distributed, indexing and meta-data search functionality. Full indexing of text, rtf, pdf, Microsoft Word and html files ensures that retrieving your content from a massive on-line archive could not be easier.

MatrixStore supports hundreds of simultaneous search per second allowing data to be shared, processed and published among authorised personnel be they internal or external to the organisation.

This functionality is built in. It does not require additional hardware or software. Searching MatrixStore is Free!

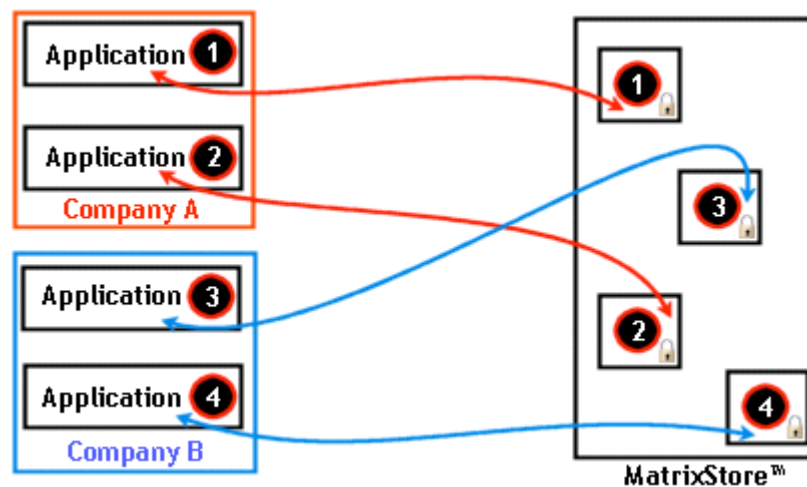


Scalability

MatrixStore has true 'plug and scale' technology; from 2 Terabytes to virtually any amount (Petabyte+). Traditional storage architectures require a large amount of planning and expertise in order to successfully grow the available capacity to the application or user. Adding capacity with MatrixStore is relatively management free, simply plug in the node, connect it to the switches and it joins the cluster, ready for use automatically. Each additional node also provides extra bandwidth for data access.

Multiple-Tenancy

It should always be the case that a storage device can be used by multiple applications and each application should be able to securely keep logically separated data vaults on that device. MatrixStore provides just that.




MatrixStore provides multi-tenancy functionality. MatrixStore provides each application with a virtual space or Vault (See [Vaults](#) for a full description) which can be expanded seamlessly without manual intervention. It is possible to set

default policies on each vault according to the storage requirements or the classification data being stored by the application.

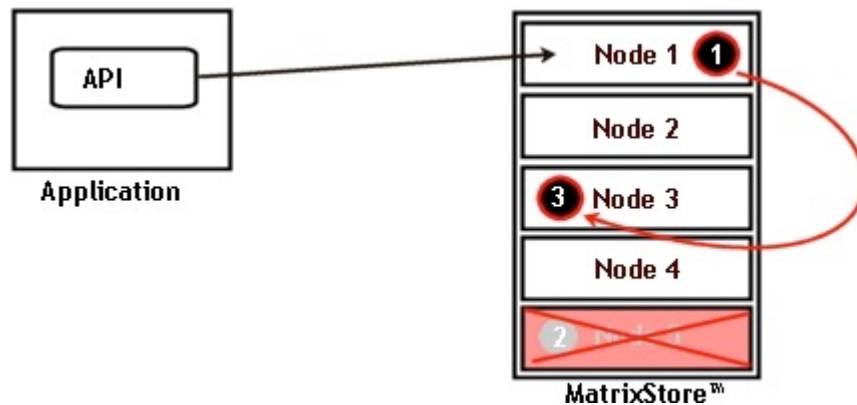
Data Protection

MatrixStore has 99.999% reliability for data availability. In a standard configuration, data is mirrored across 2 nodes. The data is written to the first node the application communicates with then synchronously copies it to another node in the cluster.



If a disk fails, then data is protected within each node, through the RAID configuration. The administrator then simply removes the bad disk and replaces it with a new disk. The RAID file system will automatically re-generate itself and return to the previous level of protection.

Nodes in a MatrixStore system can fail, MatrixStore accounts for this by monitoring each node and acting in the case of failure. When a node fails the data is protected into another node thus keeping the data both safe and available. Data on the failed node can be extracted should the disks still be intact.



Security

MatrixStore tackles data security at every critical level from read-only file-systems to auditing operations and actions on the data. All transmitted data packets are fully encrypted and authenticated. This provides protection against spoofs, and replay attacks. Because all transmitted packets are 256 bit encrypted (default) all communication is confidential.

MatrixStore even protects itself against un-authenticated nodes trying to attach themselves to the cluster.

Load Balancing

MatrixStore™ takes care of the load balancing for attaching clients and the distribution of data within the cluster. There are different modes of load balancing all of which are used depending on the current state of the cluster. There is no need to place load balancing hardware in front of a MatrixStore cluster.

SNMP and Monitoring

MatrixStore will send out status traps to an SNMP console specified (by IP address) in the MatrixStore Admin app. The traps indicate if MatrixStore requires investigation or manual intervention. The Traps sent are:

- Green (MatrixStore is operating without failure (Hardware or software))
- Amber (MatrixStore is operating but there may be a hardware failure)
- Red (MatrixStore requires urgent attention)

Statistics

Statistical information is available at MatrixStore, Vault and Node level. All the following statistics can be viewed using the MatrixStore Admin application.

MatrixStore

<i>Statistical Group</i>	<i>Values</i>
Health status	<p>Green – MatrixStore fully functional with no hardware or software failures.</p> <p>Amber – MatrixStore fully functional but will need to be serviced in order to replace failed components. Applications may still carry out all operations.</p> <p>Red – MatrixStore needs urgent Attention due to failure of components which may mean that data is no longer immediately available. Writing operations in this state are not permitted.</p>
Capacity Status	<ul style="list-style-type: none">– Free Capacity– Used Capacity– Number of Objects
Connectivity Status	<p>A list of the external servers using the MatrixStore API who are currently connected to the cluster. Details in the list include:–</p> <ul style="list-style-type: none">- IPaddress/Port of the connected server- Vault Id to which the user is connected- User Id with which the user is connecting- Creation time of connection- Last Used date on connection
Miscellaneous	<ul style="list-style-type: none">– List of Vaults created (including capacity report on the vaults)– List of Authorized Datasync Vaults

Vault

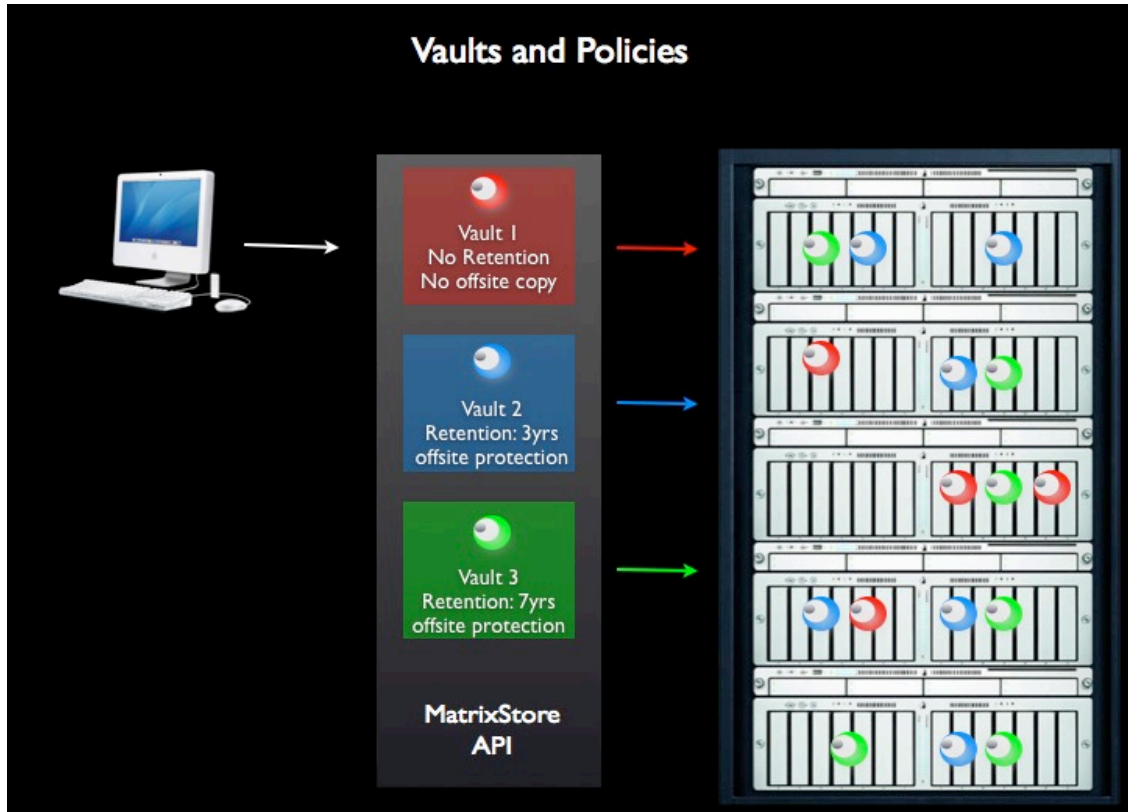
Capacity Status	<ul style="list-style-type: none"> - Number of Objects - Used Capacity <ul style="list-style-type: none"> - Total Written Ever - Total Written less Deleted Objects - Used Bandwidth <ul style="list-style-type: none"> - Total bytes written (monthly and Ever) - Total bytes read (monthly and Ever) - Monthly Capacity report (The last 13months records are kept of data in and out of the vault). Very useful for charge back within an organisation or when MatrixStore is used as part of a service offering.
-----------------	--

Node

Health status	<p>Green - the node is fully functional</p> <p>Amber - the node is functional but there may be a single component failure, i.e. one of the NICs may be out of action, data is still available.</p> <p>Red - the requires urgent attention.</p>
Uptime	Time since the node last rebooted
Version	Version of MatrixStore server software currently running on the node
Hardware/Software Warnings	<ul style="list-style-type: none"> - NIC, Power, Disk Failures are reported. - Severe Software Warnings are also reported.
Capacity Status	<ul style="list-style-type: none"> - Free Capacity - Used Capacity

Vaults

MatrixStore objects are stored in vaults. A Vault is a secure virtual storage space within that form an integral part of the MatrixStore security, compliance and multiple-tenancy functionality:-



When a vault is created it can be given a set of properties dependent on the nature of the data to be stored within it. Those properties are:

- Provisioned Capacity
- Regulation Compliance Settings
- Audit Settings
- Data Protection Settings
- Data Indexing/Content Search Setting

All data written to MatrixStore by an application will use a predefined Vault.

Provisioning

It is possible to set and update a boundary on the capacity that a vault can consume within the cluster. This is particularly useful for organisations who wish to share a MatrixStore cluster amongst departments where use of the storage needs to be governed.

Should the provisioned capacity be exceeded the users of the vault will not be able to perform any further write operations until such time as the provisioned capacity is extended or the user deletes existing data in their vault thus freeing up capacity.

It is also possible to set the provisioning to be boundless, as such a vault can grow and grow as long as there is capacity available.


Compliance

MatrixStore supports global regulatory compliance requirements by ensuring that all data stored is highly secure, authentic, available at all times and protected from unauthorised deletion. MatrixStore also provides a full audit log for actions performed against the cluster.

MatrixStore helps to support, amongst others, the:-

1. Security and Exchange Commissions rule Rule 17a.4 that aims to prevent overwriting, erasure or alteration of records.
2. HIPAA privacy ruling for Data Protection, requiring compliant backup methodologies to ensure the security and confidentiality of patient records.
3. The Sarbanes-Oxley Act of 2002 protecting investors by improving the accuracy and reliability of corporate disclosures. The Act amends mail and wire fraud infractions with harsher punishments and imposes fines and prison sentences of up to 20 years for anyone who knowingly alters or destroys a record or document with the intent to obstruct an investigation.

The MatrixStore solution is only one piece of a compliance strategy. To ensure that your company goes about its business in a compliant manner the correct processes, training and software applications must also be put into place.



Vaults need to be assigned a compliance policy, this is a one time choice as once the type of policy has been set it cannot be changed. This policy can be any of the following:-

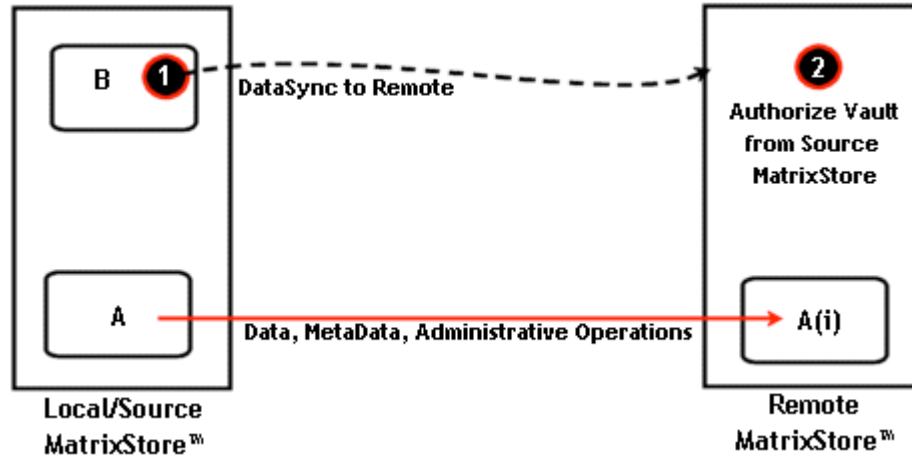
- None (data can be deleted at any time)
- Extendible (the compliance threshold can be extended after it has been initially set)
- Reduceable (the compliance threshold can be reduced after it has been initially set)
- Both (the compliance threshold can be reduced and extended after it has been initially set)

Whilst the compliance threshold is still valid data cannot be deleted or changed. When the compliance threshold has expired the application will be able to perform these mutating operations.

The choice of Compliance Policy should be governed by the classification or type of data to be stored in that vault in accordance with any internal or legislative requirements.

Data Protection

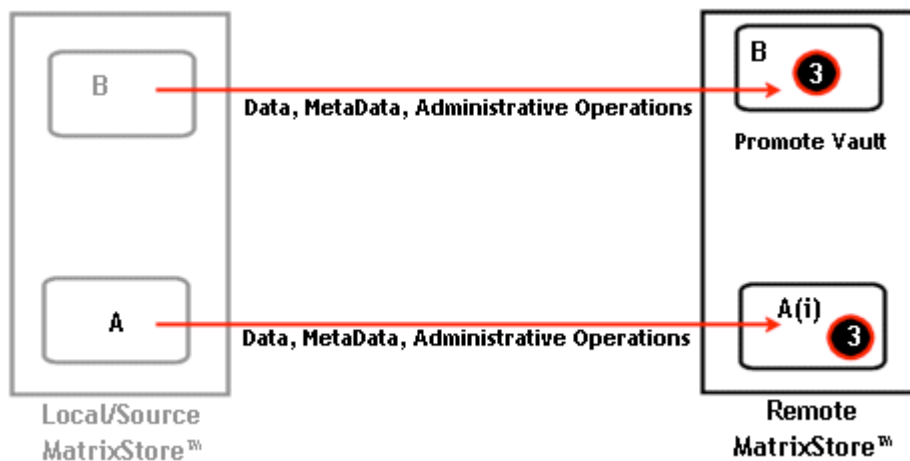
MatrixStore replication is an **Asynchronous** cloning mechanism which ensures that any actions performed on a local vault are also carried out on the remote vault, actions include data and administration operations. All replication communications are performed over secure, encrypted channels.



The remote vault is purely a remote backup, as long as it is tagged as being a 'replica vault' users will not have the same capabilities on the remote store as they do on the local source vault. The remote replica vault is Read-Only thus allowing the application to perform a read fail-over should the source vault not be available temporarily.

The semantics of replication are :-

1. Create a vault on the source MatrixStore cluster.
2. On the remote cluster you need to authorise a source vault to replicate. This is a safety mechanism which controls where the remote MatrixStore receives its data from



3. When the source MatrixStore or vault becomes permanently unavailable the remote vault will need to be promoted to be the new source cluster. Once promoted users of the vault can connect directly to the store and perform all data and administration operations on this store.

Audit Trails

Every operation on MatrixStore or on a vault is audited. Auditing access to objects is not only crucial in order to comply with global regulations but can be very useful when MatrixStore is deployed in industries that care about logging the distribution of valued digital assets, organisations such as broadcasters and record companies who require strict controls over digital rights.

The audits are split into two functional areas; **administrative** operations and **data** operations.

Administrative Operations	Administrative operations include: <ul style="list-style-type: none">● Changes to vault settings● Creation, modification, deletion of users● Failed login attempts
Data Operations	Data operations include: <ul style="list-style-type: none">● Object creation● Object deletion● Update of object attributes● Object Access (optional audit of reads)

The data and administration audit logs also contain information such as time-stamps, vault/user Ids, the target of the action and the action performed. The audit information is kept secure on each node and is protected at the same high level as is the customers data. At present it is possible to retrieve the whole audit trail for a given Vault or view the trail for a given time period using the MatrixStore Administration Tool.

User Capabilities

It is possible to set capabilities at vault level. All users of the vault can only perform operations based on capabilities that are the same as or, are a subset of those belonging to the vault. The capabilities are:-

- Write
- Read
- Search
- Delete

It is possible to change the capabilities of a vault at any time which will have could have a direct on users of the vault, i.e., it is possible lock down a vault by disabling read/write operations at vault level thus stopping all users of a vault from performing those operations.

Protection Algorithms

MatrixStore validates the authenticity of data written by utilising hashing algorithms. (Adler32 and MD5)

When data is written to MatrixStore the protection algorithm is used to verify that the data is authentic once it has been committed to the disk. The hash is generated on both the initial node the write requests commences on and the node to which a second copy of the data is relayed to. If the generated hash values on both nodes do not match then the write is deemed to have failed.

The generated hash is also stored as an attribute on the object itself for future validation.

Data Operations


All operations are performed using the MatrixStore API. See the API section for more details and a link to the full MatrixStore API documentation.

Write

It is possible using the MatrixStore API to write data of any type and of any size (depending on the amount of space left on two volumes within the cluster of course!) in size for a single item. It is possible for meta-data to be associated with the data at object creation. The MatrixStore API and MatrixStore server support concurrent access to a single or multiple vaults so data can be written in parallel from multiple clients to the same or different vaults. Please refer to the API Reference and User Guides detailed specifications and examples.

The following is an outline of the write path:-

1. Application connects to MatrixStore using API.
2. Application requests to perform a write operation.
3. MatrixStore chooses the best location to store the object.
4. MatrixStore returns a list containing optimal locations to write the data.
5. The API connects to the first node in the list and starts to pump the data.
6. The data is relayed to the second node in the list.
7. Once the data is on both nodes a hash of the data is taken on each node and compared to ensure that no corruption occurred during the transfer.
8. MatrixStore returns a globally unique object Id to the API and thus the application.



Read

The MatrixStore API and MatrixStore server support:-

- **Concurrent access** to a single object and thus many clients can read the same piece of data at the same time without being queued. Please refer to the API Reference and User Guides detailed specifications and examples.
- **Partial Object Retrieval** such that a section of a large object can be retrieved given the start byte position is provided.
- Multi-Threaded access to the server such that many objects can be retrieved simultaneously.

The following is an outline of the read path:-

- Application connects to MatrixStore using API
- Application requests to perform a read operation passing in the globally unique object Id and optionally a start byte offset
- MatrixStore locates the nodes that the object resides on.
- MatrixStore returns the node locations to the API in a list. The first node in the list being the least loaded and thus most appropriate node to serve the data.
- The API retrieves the data, building up a local copy in the location specified by the application.

Search

MatrixStore supports the following modes of searching for meta-data associated with objects stored in the cluster.

- Simple Search Term
 - o meta-data Key/Value pair
- Complex Search Term
 - o Series of AND/OR over multiple simple search terms

MatrixStore supports concurrent search operations which can each return hundreds of search results per second.

Delete

Deleting an object is possible should it not be under a current and valid compliance threshold. If the object has already been deleted the application will be informed that the object has been deleted and that a 'tombstone' exists detailing how and when it was deleted.

The following is an outline of the delete path:-

- Application connects to MatrixStore using API
- Application requests to perform a delete operation passing in the globally unique object Id
- MatrixStore locates the nodes that the object resides on and if not protected by a current and valid compliance threshold the data is deleted.
- A Tombstone is created to leave a trail of when and by which user the object was deleted.

Update Attributes

It is possible using the MatrixStore API to update the meta-data associated with data stored in MatrixStore. This is a useful feature for tracking use of an object throughout its lifetime in the archive or it could be that new data is available on a historical piece of data stored in the archive and the meta-data needs to be updated to reflect that.

MatrixStore Interfaces

MatrixStore API

The MatrixStore API is available in both C and Java. A typical integration with the API should take less than a week development time plus any testing cycles. Please refer to the MatrixStore API User and Reference Guides for more detailed information.

MatrixStore Management API

All administrative operations are carried out over a secure 256bit encrypted link using the MatrixStore management API. At present the management API is not open to develop against.

DropSpot

DropSpot is a powerful, secure and flexible archive and workgroup desktop application which interfaces with MatrixStore API. With DropSpot it is possible to perform the following operations:-

- Archive/Write to MatrixStore
- Read from MatrixStore
- Search MatrixStore
- Delete from MatrixStore
- Share data amongst an authenticated workgroup



It is possible to run MatrixStore in both GUI and console mode which allows administrators the flexibility to either drag and drop archive jobs or run overnight utilities.

DropSpot is a Java client which runs on the Mac OS, Windows 2000/XP and Linux platforms.

Please refer to the DropSpot user guide for more information.