



# Email Archiving Success: Message Warehousing on the Mimecast Storage Grid

Centralized email archiving for the enterprise

## Message Warehousing on the Mimecast Storage Grid

### Centralized email data management for the enterprise

“ Creating a secure central repository of all business email data can be an expensive and technically complex undertaking. ”

“ In-house email archiving projects consume a disproportionate percentage of budgets relative to other IT initiatives. ”

“ Data volumes simply grow too large too quickly to store and secure effectively within budget, time and technology constraints. ”

Today centralized email data management is a critical requirement for every business. Without it companies lose control over their important electronic business records and valuable information, and cannot effectively meet operational and legislative requirements.

Organizations that leave the management of email data in the hands of their end users lack the control needed to guarantee that their data is secure, discoverable and available when it is needed.

Creating a secure central repository of all business email data can be an expensive and technically complex undertaking. The larger an organization, the more mail servers and locations need to be incorporated, the more challenging and indeed the more important it is to have centralized email data management.

Various email archiving products are available on the market, each requiring the purchase and deployment of expensive redundant hardware, software and complex integration on local networks before a benefit can be realized. In-house email archiving projects tend to consume a disproportionate percentage of budgets relative to other IT initiatives. This results from their complexity, the requirements for offsite back-up and the growing and unpredictable data volumes involved.

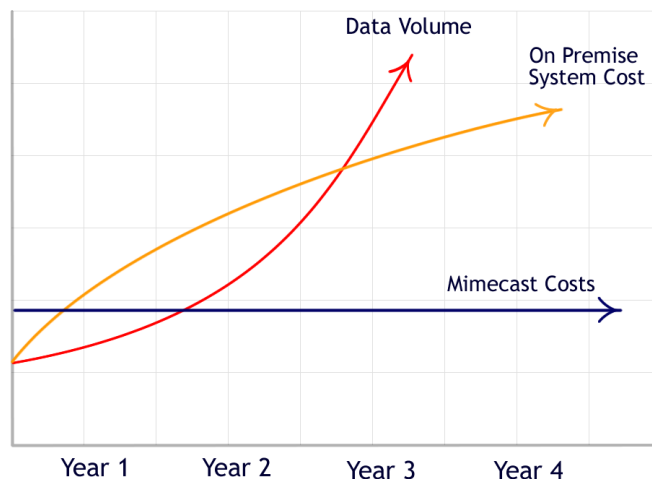


Fig. 1: Mimecast vs. On-Premise Archive: Cost & Data volume comparison

Most IT managers are new to the responsibility of maintaining the availability of email data for multi-year periods. A significant amount of new email data is created and received every day and each day marks the beginning of a new multi-year data retention period.

“ A simple operational requirement quickly turns into a significant long term financial and technical burden. ”

Business Continuity planning dictates that this data must be stored on resilient hardware and be backed up, ideally in real-time, to one or more offsite locations. As data volumes grow, traditional backup and recovery strategies or disk to disk replication approaches become less effective.

Data volumes simply grow too large too quickly to store and secure effectively within budget, time and technology constraints. This information is also expected to be online, searchable and instantly retrievable by the business at all times.

Attempting to meet these requirements with in-house systems drives significant capital and operating expenses into a business and often result in a suboptimal solution. Additional costs are to be borne year over year as the business extends and patches what soon becomes a legacy data repository.

A real challenge faced is when onsite email archiving systems do not scale at a software level and cannot perform adequately when required to handle large and rapidly growing volumes of data. With message volumes on the rise and average file attachment sizes increasing steadily, correctly sizing and scaling an in-house email data management system inevitably involves guesswork. Mistakes can be costly as large volumes of data inevitably encounter a performance degradation threshold.

Changing platforms will involve significant expense as well as downtime and potential data loss. Besides software scalability, hardware and backup facilities also need to scale and providing these facilities can involve substantial recurring costs. A seemingly simple operational requirement can quickly turn into a significant long term financial and technical burden.

Because effective email data management is vital to every business, IT departments need a practical strategy that is low risk, future proof, immediately affordable, easy to implement and manage, and critically, has a predictable cost profile over the many years of email data management that lie ahead.

#### **The solution:**

#### **Message Warehousing on the Mimecast Storage Grid**

Mimecast provides everything an enterprise email data management strategy needs in a ready-to-use online best practice configuration.

The backbone of Mimecast's data retention service is Mimecast Stor4 - a geographically diverse highly available, secure and resilient storage and information management grid. Mimecast Stor4 is a unique software system designed by Mimecast to meet the new business challenge of providing high performance high availability data retention services to thousands of customers simultaneously, on an extremely cost effective Internet based system.

Mimecast Stor4 is fully integrated into the Mimecast platform so that clients can benefit from a completely unified and holistic email management service. The Mimecast Stor4 grid is served from a geographically diverse best of breed storage hardware and high performance 64bit computing infrastructure.

“ Effective email data management is vital to every business. ”

“ Mimecast provides everything an enterprise email data management strategy needs in a ready-to-use online best practice configuration. ”

“ Mimecast offers the industry's most sophisticated multi-tenant always on email solution. ”

“ All data is immediately written to at least three of its geographically diverse storage nodes, each housed in secure interconnected data centres. ”

This infrastructure is housed in highly secure data centres behind layers of advanced security technology. Mimecast offers the industry's most sophisticated multi-tenant online email storage solution with unmatched levels of search performance, data security, availability and resilience.

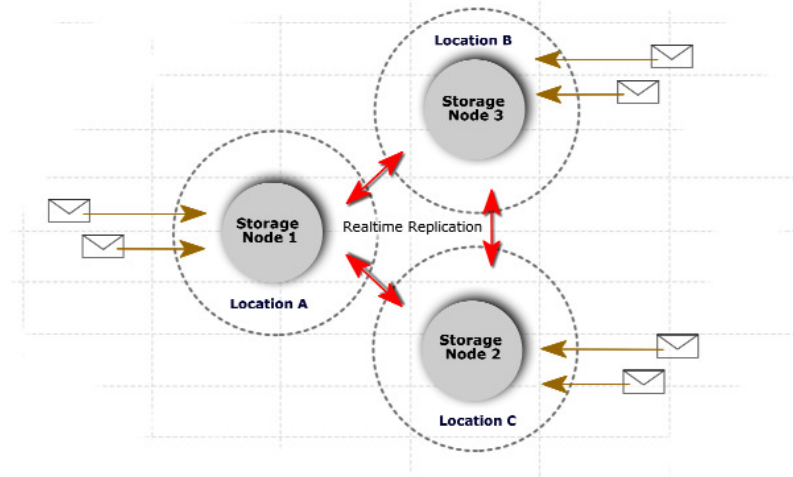


Fig. 2: Stor4 storage nodes and replication paths

Mimecast provides the software application services as well as the physical network, computing and storage infrastructure required to deliver comprehensive unified email data management. The service incorporates real-time email continuity services, long term email retention, compliance, discovery and information management tools, comprehensive gateway based email security and policy controls and the most sophisticated online MTA and enterprise intra-organization routing functions available in the market.

With this technology Mimecast is responsible for the safe delivery, storage, and management of hundreds of millions of messages for thousands of businesses worldwide every day.

**Critical capabilities to deliver against your enterprise email data management strategy:**

**Triple Resilience** - As data is received by Mimecast it is dynamically written to the Mimecast Stor4 grid. The grid software ensures that all data is immediately written to at least three of its geographically diverse storage nodes, each housed in secure interconnected data centres. Each node has built-in redundancy to protect against individual device and media failure, as well as self healing connectivity to other nodes to ensure completeness of resilience.

**Dynamic Data Protection** - Each node on the Stor4 grid is aware of its remote storage partners for all segments of data and is aware of its own role in protecting every piece of data within each segment. Should a node suffer from a partial or complete loss of data caused as a result of media failure or a catastrophic local disaster, normal system operations will continue unaffected. Once new media or a new node is allocated for redundant recovery, the failed node's remote partners will automatically rebuild the missing data and ensure that the grid maintains triple redundancy.

“ Mimecast is the only provider with the technology to offer massively parallel search aggregation delivering sub-second full text search results. ”

“ All access to your data by your administrators is logged in a comprehensive access audit trail. ”

“ Mimecast's unique software architecture that ensures the continuous availability of data from multiple nodes on the grid at all times. ”

**Advanced Data Integrity** - The Mimecast Stor4 software performs constant data integrity checks to ensure that the data written to it is valid and complete at all times. Should any element of the data become corrupt, the system will detect and repair itself automatically from the other valid copies of the data available on the grid. This prevents any wide spread corruption of data or an error outbreak condition as seen with other 'classic replication model' devices, where errors and corruptions are blindly replicated across the dataset.

**Rapid Data Acceptance** - Given the volume of email data an enterprise can produce across its many mail servers on a daily or even an hourly basis, a centralized repository needs to be capable of reliably receiving, indexing, and securing a constant stream of new content.

Mimecast Stor4 has highly optimized software processes to meet this challenge and balances loads across powerful multi-core node controllers with the raw processing power to accept substantial streams of data in real-time. A clients email data comes from many different mail servers both inside and outside the business. Therefore Mimecast accepts external emails via the SMTP channel and internal mail from your various internal mail servers via a secure and native Microsoft Exchange Journal Hosting connector.

**Always On** - Mimecast offers you guaranteed access to your data 24/7 for the full duration of your contracted retention period. This guarantee is only possible due to Mimecast's unique software architecture that ensures the continuous availability of data from multiple nodes on the grid at all times. This is important as no company can predict when historic email data will be required by a user or to meet a strategic business requirement. Mimecast therefore keeps all of your data available to you at all times without compromise.

**Search Scalability** - Without significant search scalability, a centralized enterprise email data infrastructure will become slow and ultimately unusable. Delivering sub-second full text search and retrieval across large data volumes requires different software architectures and infrastructure than standard email archiving platforms, and these are generally too expensive for normal organizations. Mimecast is the only provider with the technology to offer massively parallel search aggregation capabilities to fit the budget. These capabilities are essential for keeping search and retrieval performance acceptable across an enterprise's full email data volume.

“ Given the volume of email data an enterprise can produce across its many mail servers, a centralized repository needs to be capable of reliably receiving, indexing, and securing a constant stream of new content.”

“ Mimecast offers rich message mining capabilities that make tasks that seemed impossible in the past happen in milliseconds.”

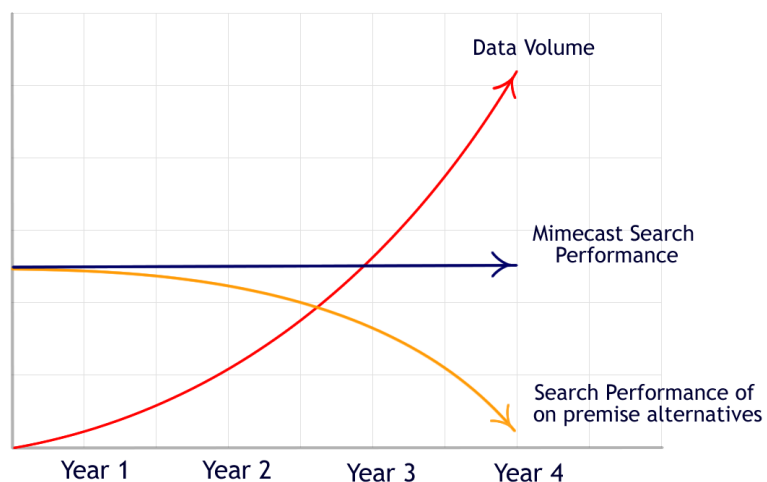


Fig. 3: Mimecast vs. On-Premise Archive: Search Performance

**Tamper Proof, Auditable and Secure** - Access to your data is only possible by your authorized administrators via a secure authenticated log-in to your Mimecast account. Mimecast supports authentication against your own Microsoft Active Directory as well as a local security database. All access to your data by your administrators is logged in a comprehensive access audit trail. Each file is encrypted uniquely for each client and receives a unique checksum on entry. Mimecast can therefore certify that client's data has not been modified since it entered the grid.

**Completeness of data** - An email data retention strategy is incomplete unless all the information pertaining to a message, its various versions, components, policies enforced upon it and forensic trace information is retained too. Forensic trace information such as proof of delivery information which is usually only held temporarily at the Internet email gateway level is critical to being able to prove data validity.

Mimecast retains all the key information such as which mail server the message was accepted from and delivered to, the greeting received by the respective servers, the data transmission start and end timestamps, and the final acknowledgments issued from and to these servers. Without this information you may not be able to conclusively prove that a message was delivered by you to your intended recipient, or indeed that a message was received from who you purport it to have come from.

Mimecast is the only email data management platform that records the full trace data-set. Other on-premise archiving tools are disconnected from the actual receipt and delivery actions performed at the gateway and as such lack the completeness and integrity of this vitally important data.

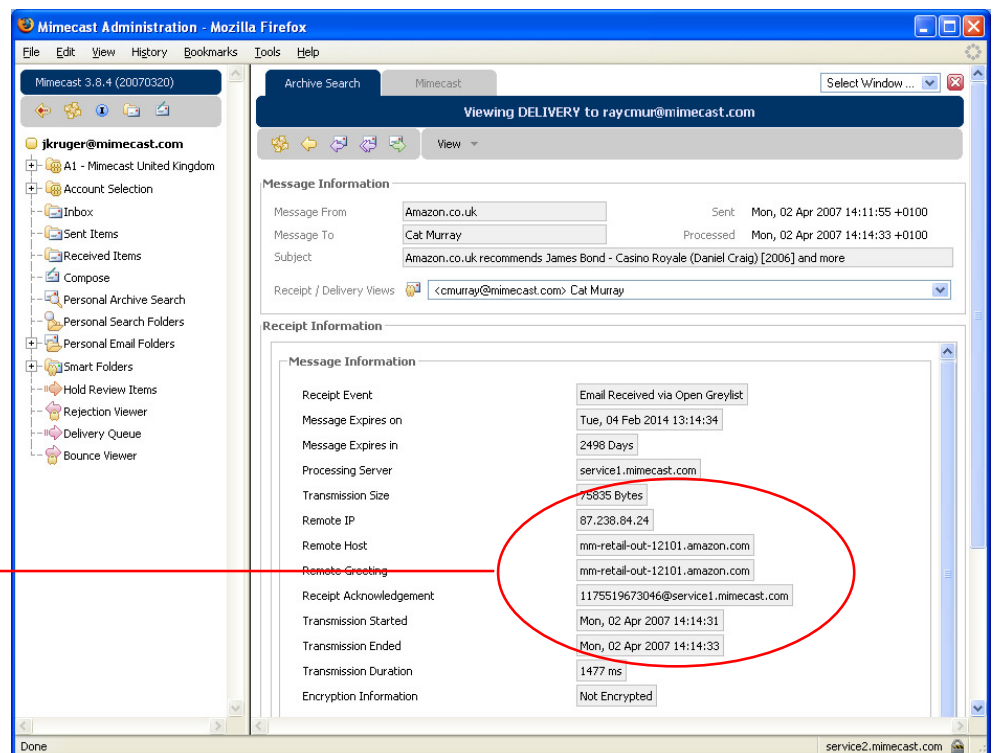


Fig. 4: Forensic track and trace

Proof of delivery information

**Flexible information access** - Providing useful access to vast repositories of email data can be challenging. Mimecast provides flexible tools to enable you to meet the demands for information that your business will have. For example; saved searches, smart folders with permission based views, export as PST options, personal archive searches, and content or document specific searches. Mimecast offers rich message mining capabilities that make tasks that seemed impossible in the past happen in seconds.

**Single view of global email data**

Delivering a single online archive and searchable view across an enterprise's entire global email dataset is not trivial. Mimecast delivers this seamlessly by rapidly aggregating and indexing data from any number of mail servers on your networks around the world. This removes the need for multiple fragmented local email archive systems and complex discovery processes. Mimecast Stor4 delivers the storage and search power to offload the intense demands of global email data management from your in-house resources.

Mimecast provides your business with a powerful corporate memory that can be accessed quickly, easily and securely to meet today and tomorrow's business demands.

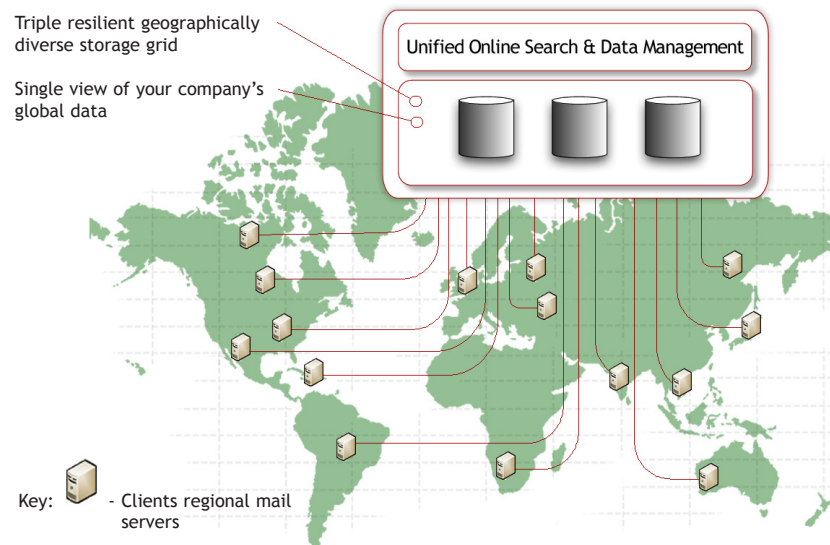


Fig. 5: Global email data aggregation

**About Mimecast**

Mimecast is the leading innovator in the Software as a Service business email management market. The company provides best practice messaging security, storage, continuity, business information management and policy control to clients around the world via a unified internet based architecture. Mimecast was founded in 2002, and is headquartered in the UK.