

Cryoserver™

Forensic Compliance System

Technical Overview White Paper

V4.2



1. EXECUTIVE OVERVIEW

1.1 Purpose

This document describes the architecture of Cryoserver and explains how the Cryoserver modules work together.

1.2 Introduction



The New Cryoserver Appliance repositions the Cryoserver forensic email compliance solution from the pure security sector into email/IM security and storage management. Cryoserver eliminates the need for email backup, DR and quotas while providing unprecedented compliance & lifecycle management. Simple to use, installed in minutes, providing email protection forever.

Cryoserver is considered one of the leading Forensic Compliance Systems. It takes an audit copy of electronic communication sent to, from and around an organization in real-time. Data is held within the Cryoserver Appliance in an encrypted, tamper-evident environment that can be searched extremely quickly. This is suitable for use in legal/regulatory compliance, producing data in court with high evidential weight and for speedy resolution of disputes and other situations that require or demand electronic messaging as evidence. The Cryoserver repository of messages is designed to prevent arbitrary and unauthorized deletion or alteration.

Additional technical information is provided to help the reader assess the infrastructure impact and requirements of implementing Cryoserver.

1.3 Overview

Cryoserver accepts email messages, either by network interception or more typically using journaling, POP or IMAP collector, from a variety of Email messaging platforms. These include Microsoft Exchange, IBM Lotus Notes, Novell GroupWise, Sun Java Enterprise System (JES), Scalix, Fujitsu Teamware and others.

Cryoserver captures Instant Messages using network interception.

Cryoserver copies messages for storage and indexing. A search engine is provided that can query large volumes of archived data by content, including any attachments and metadata. Cryoserver is a distributed system and designed to be scalable and robust.



2. ARCHITECTURE

2.1 Design Goals

Cryoserver has been designed to aid organizations requiring a totally dependable system for storing electronic messages for later retrieval. We believe that Cryoserver satisfies the requirements for regulatory authorities and is able to provide a trail of data with high evidential weight for use in a court of law.

Distributed	A distributed architecture enables the system to be configured to avoid processing, storage and network traffic bottlenecks.
Appliance independence	Cryoserver is delivered as independent pre-built Appliance solution. The solution has also been certified on IBM, Sun, HP and Dell servers.
Scalable	Cryoserver's Appliance architecture means that it can be scaled to meet existing and future requirements simply by adding more appliances as necessary, without impacting the existing storage and indexing configuration.
Open architecture	Cryoserver's open architecture means that IT professionals can be satisfied that it will meet required performance, scalability, robustness and security standards. It also means that Cryoserver can be integrated with third-party email, storage and indexing systems.
Robust	Full use is made of hardware and operating system level redundancy capabilities, integrating readily with an organization's business continuity and disaster recovery programmes.

2.2 Key User Features

Cryoserver has been designed to be intuitive and easy to use.

Searching	Relevance-ranked reporting of messages that satisfy a user's search criteria. Capability to sort the reports according to user-specified requirements (size, date etc.) Dynamic links are provided directly to summarised email (and attachments).
Flexibility of Searching	Simple search on content and metadata plus sophisticated searching utilising stemming, sound-alike and proximity.
Email recovery	Retrieval of messages by individual users with the 'forward to inbox' facility.
Secure access	The front-end is delivered seamlessly through Outlook or over HTTPS and users are authenticated against LDAP or Active Directory.
Remote access	Remote access through any appropriate browser system (subject to security constraints).
Credibility	As Cryoserver retains a copy of messages and their associated metadata, in real time, messages retrieved demonstrate a full chain of communications and be more credible than data retrieved from servers or back-up tapes (which may not show a full record as



deletions/alterations are possible).

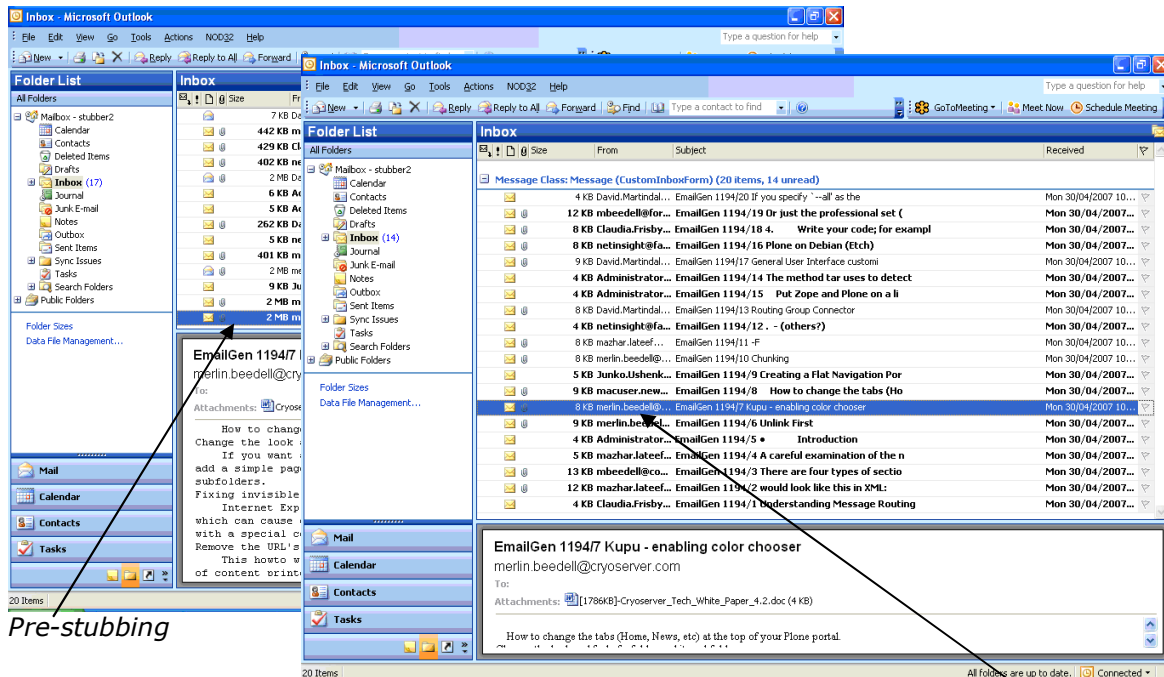
3. STORAGE MANAGEMENT & STUBBING

Storage Management and Stubbing:

The Cryoserver architecture also supports storage management features with the 'mailbox management' or stubbing facility. Companies, if they choose to, can eliminate the need for quotas and message size restrictions and give users a mailbox of virtually unlimited size while also controlling message store growth. Exchange store management significantly reduces the size of Microsoft Exchange stores, further reducing e-mail storage by as much as 80 per cent¹. Furthermore Administrators can significantly reduce the time spent dealing with mailbox housekeeping. Combined with other features of the Cryoserver solution, the need for Exchange Server backup operations is significantly reduced (and can be eliminated).

Automated policies on the Exchange server replace each email message with a small kilobyte stub (referrer) which points to the email in the Cryoserver repository instead. This gives immediate benefits back to an organisation, as less expensive primary disk is utilized on the Exchange server, an application server. Reduction in storage volumes being used gives a boost in performance to Exchange and resulting in end user productivity. Original emails are replaced, for example a 10MB email can be replaced by a 5kb stub in user's mailboxes which opens the email from the Cryoserver repository instead. To the end user it is a transparent process and they would not even know that the mail is not being viewed from Exchange.

This policy can be based on date, size of email or mailbox name. You can even have emails arrive in users Outlook mailboxes already stubbed, stubbing in real time, which can lead to multiple benefits including but not limited to less bandwidth utilised; even remote workers synchronizing with Exchange server can be faster. All of this is transparent to the end user.



Pre-stubbing

¹ The actual amount of storage saving is dependent on the stubbing rules defined. The more lower the stubbing rule latency (i.e. the shorter the time emails are retained) the greater the storage savings.



Post Stubbing

The mail residing on Exchange and Outlook is 256 times smaller. The Cryoserver stub changes the mail from 2MB to 8KB. Each attachment is replaced by a 4K stub or referrer. This process is transparent to the end user.



4. CRYOSERVER ADMINISTRATION

Within Cryoserver there are four types of users: basic, privileged, administrator and superuser. Cryoserver administration is carried out by administrators and superusers.

4.1 Administrators

Administrators are responsible for creating and maintaining Cryoserver specific accounts, which are normally restricted to a small number of privileged and administrative accounts. All actions are logged for audit purposes; administrators have no access to the email repository.

When an administrator logs in they are directed to the User Administration section of Cryoserver, which is shown in the image below:

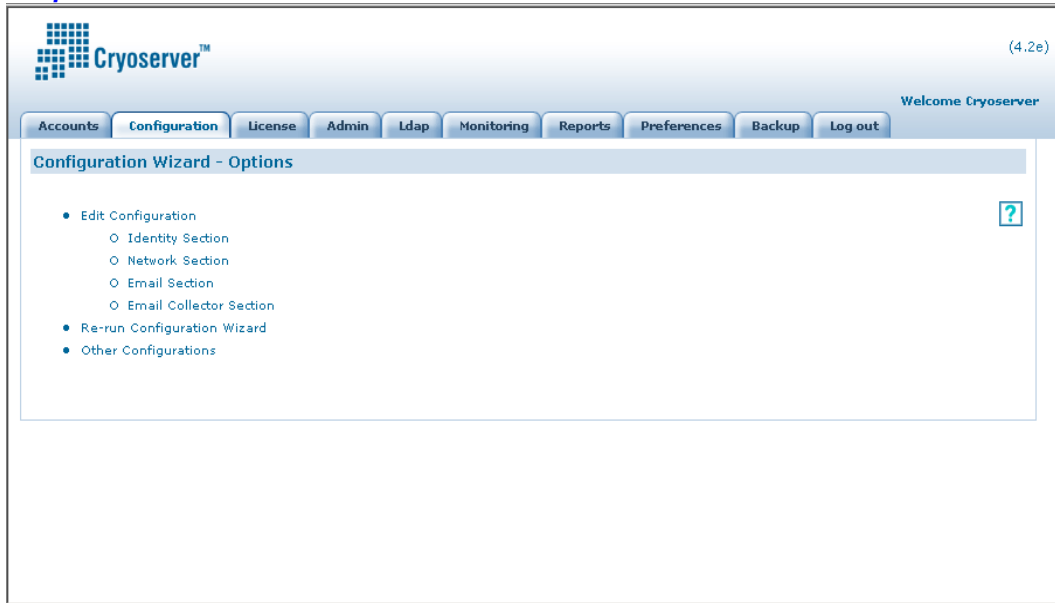
Administrative user interface

The screenshot displays the Cryoserver administrative interface. At the top left is the Cryoserver logo and the version number (4.2e). A navigation bar contains tabs for Accounts, Configuration, License, Admin, Ldap, Monitoring, Reports, Preferences, Backup, and Log out. The 'Accounts' tab is active, showing a list of 'Existing Accounts' on the left: amit.kumar, cryoserver_admin, and priv. The main area is titled 'Account Details' and contains several input fields: Username, First Name, Last Name, Admin Level (set to Basic), Account Status (set to Active), Primary email address, and Secondary email addresses. There are 'Add' and 'Remove' buttons for the secondary email addresses. At the bottom, there are fields for 'Last log-in date' and 'Account creation date'. 'Create User' and 'Cancel' buttons are located on the right side of the form.

4.2 Superusers

Superusers access a separate administration area to modify the Cryoserver configuration and monitor the Cryoserver system. The superuser interface is shown below.

Super user interface



Company Details include, Configuration, License Key, Reports, Data Guardian, retention period, SMTP address, LDAP configuration, email format and data guardians.

The Monitoring screen (shown below) is important as it provides the ability to view the status of all of the Cryoserver modules (nodes) to be viewed, even if they are installed on different servers in a distributed configuration.

The image below shows the status of (in order), the Storage Director (SD), the Search Engine (SE), the Spool Manager (SP), the Index Manager (IM) and the Storage Manager (SM). Where mirrored pairs of managers are used for resilience, the Partner column identifies the partner of each mirrored node. For nodes that use disk space the Capacity column shows the percentage and actual amounts of disk space currently used by the node.



Monitoring page of the superuser interface

Cryoserver™ (4.2e) Welcome Cryoserver

Accounts Configuration License Admin Ldap **Monitoring** Reports Preferences Backup Log out

Refresh Index Now ?

Site: production

Node	Host	Partner	State	Information
uk-ln-sd-001	ps029	-	Online	Wed Jun 13 11:39:29 GMT+05:30 2007
uk-ln-se-001	ps029	-	Online	-
uk-ln-sp-001	ps029	-	Online	spool size = 0, error queue size = 0, processed = 2 agent 0: none, 0 b (20 minutes) agent 1: none, 0 b (20 minutes)
uk-ln-sn-001	ps029	-	Online	0.00% used (2.3 Mb/51200 Mb) . Merging index for [cryoserv/email]
uk-ln-nt-001	ps029	-	Online	-
uk-ln-ds-001	ps029	-	Online	-
uk-ln-rm-001	ps029	-	Online	-

Repository Status

Company	Document Count	Error Queue Size
cryoserv	2	-
unknown-account	-	-
zerosize-account	-	-

Build Details

Build Number	40
Build Date	08-Jun-2007 15:08 GMT+05:30
Java Version	1.6.0_01-b06

List of currently logged in Users

User Name	Company	Admin Level	Connected From
priv	cryoserv	Privileged	127.0.0.1
cryoserver_admin	cryoserv	Administrator	127.0.0.1
priv	cryoserv	Privileged	127.0.0.1

All activity performed by the superuser is logged in Cryoserver and the audit log is emailed to the Data Guardian(s), who watch over the system.



5. SEARCH AND RETRIEVAL

The Cryoserver email repository can be searched by basic and privileged users only.

5.1 Basic Users

Basic users may use Cryoserver's powerful email search facilities to access their own repository of emails. They normally access Cryoserver using their normal network login using LDAP / Active Directory authentication, although it is possible for administrators to create basic user accounts within Cryoserver if necessary.

5.2 Privileged Users

Privileged users have the ability to search emails throughout the entire repository. This level of access is intended for a few trusted individuals (such as a Compliance Officer, HR Manager or Data Protection Official) within an organization and should be considered in co-ordination with privacy regulations, both corporate and legislative. Privileged users are required to state their reasons for searching. All searches they carry out are logged, and session transcripts are automatically stored in encrypted format in Cryoserver as well as being sent by email to nominated Data Guardians who have the responsibility for checking that searches are in accordance with the stated reason, corporate policy and regulatory requirements.

5.3 Search Interface

The standard search interface is shown below. Users can search for emails by specifying any of the search criteria. Searches can be refined where necessary by adding further search criteria and reissuing the search. Cryoserver's speedy search capability ensures that users are able to quickly find the messages they are looking for.

Basic search screen

The screenshot shows the Cryoserver search interface. At the top left is the Cryoserver logo and the version number (4.2e). On the right, it says "Welcome Robin". Below the logo is a navigation bar with buttons for "Search", "Adv Search", "Folders", "Preferences", and "Log out". The main search area contains the following fields:

- Search document types: Email (dropdown)
- Period: Date Quick Select (dropdown)
- Start Date: 29 Oct 2006 (with navigation buttons: <Y, <M, <D, D>, M>, Y>)
- End Date: 29 Apr 2007 (with navigation buttons: <Y, <M, <D, D>, M>, Y>)
- Keywords: (text input)
- From: (text input)
- To: (text input)
- Attachment Names: (text input)
- Attachment Keywords: (text input)
- Parties: (text input)

A "Search" button is located at the bottom of the form. A help icon (?) is visible in the top right corner of the search area.



The **advanced search** provides more sophisticated searching capabilities, rather similar to what you might expect from a document management system. For example, you can search for combinations of words, for “stems” (words beginning with the same characters), and for words sounding similar.

Advanced search screen

Search document types: Email

Period: Date Quick Select

Start Date:

End Date: 29 Apr 2007

From:

To: joy remuzzi

Apply search to:

- Messages
- Attachments
- Messages and Attachments
- Headers

Include documents which contain

	Keyword	No Modifier	Spelling stems	Sounds similar	Group these words within...
<input checked="" type="radio"/> All of these words	<input type="text"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>
<input type="radio"/> Any of these words	<input type="text"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>
	<input type="text"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>

Exclude documents which contain

	Keyword	No Modifier	Spelling stems	Sounds similar	Group these words within...
<input checked="" type="radio"/> All of these words	<input type="text"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>
<input type="radio"/> Any of these words	<input type="text"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>
	<input type="text"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>

Search



6. SECURITY

6.1 Introduction

In a distributed Cryoserver Appliance configuration consisting of multiple servers, the Cryoserver modules communicate with each other using RMI over SSL. Users access the search engine with HTTP over SSL.

Access to the Cryoserver queries and administration is monitored and restricted by a User ID and password login. Security can be enhanced by using two-factor authentication using the RSA ACE Server and SecurID token.

6.2 Message Encryption

All messages are encrypted using Advanced Encryption Standard (AES-128) as the encryption algorithm which offers high levels of data protection before being committed to long-term storage. To put AES into context against, for example, DES, National Institute of Standards and Technology (NIST) have estimated the time to crack a 128-bit key (assuming a machine could crack a DES key in 1 second): is a 149 trillion years.

Cryoserver uses the Java Cryptography Extension (JCE), which supports pluggable cryptography modules. It is therefore possible to support a number of different cryptographic algorithms to meet the most stringent security requirements if required.

6.3 Message Digest

Before a message is committed to long-term storage an MD5 digest based on its content is computed and stored with the message. When a message is retrieved, a new MD5 digest is computed and compared to that stored with the original message. This allows the system to detect whether the message has been tampered with since it was stored. The user interface reports the results of comparing these message digests.

6.4 Remote Method Invocation

Remote Method Invocation (RMI) is Sun's architecture for distributed Java applications. Even without encryption RMI has a reputation as a secure protocol since the bulk of the traffic consists of serialised objects that are the parameters or results of remote method calls, making reconstruction and interpretation of an intercepted conversation extremely difficult.

6.5 Secure Sockets Layer

Secure Sockets Layer (SSL) is an open standard for providing a communication link for client/server applications that prevents eavesdropping, tampering or forgery. SSL is widely accepted as the standard for secure communication over the Internet and is relied on for e-commerce and other security-critical applications.

6.6 Secure Authentication

Users connect to Cryoserver over HTTPS using a standard web browser. Cryoserver is usually configured to authenticate basic users against an LDAP-enabled directory, which avoids the need to create additional user lists within Cryoserver. Administrative and privileged users are managed within Cryoserver's built-in authentication system

The authentication modules of Cryoserver are extensible, allowing for RSA SecurID challenge-response authentication.



7. IMPLEMENTATION CONSIDERATIONS

7.1 Platforms

The standard Cryoserver modules are written in Java and have been certified on Solaris and Linux operating systems. Modules depending on third-party components such as collectors for proprietary email servers rely on platform-dependent APIs.

7.2 Firewalls

Part of a Cryoserver implementation project involves documenting the relevant network infrastructure to identify where internal firewalls may need to be reconfigured to allow Cryoserver network traffic. The ports used by Cryoserver for inter-module communication can be configured where necessary to conform to local security/infrastructure policy.

7.3 Redundancy and Reliability

In order to provide a Cryoserver Appliance solution with the highest availability and to provide resilience to a site disaster, it is recommended that Cryoserver Appliance is deployed in a mirrored configuration. In this configuration, messages received by the master Cryoserver system are immediately copied to the slave Cryoserver system. The messages are processed in parallel by the respective Storage Manager(s) and Index Manager(s) and the systems are maintained in an identical state at all times. It may not be necessary to perform tape backups when a mirrored Cryoserver configuration is deployed, which provides the very highest levels of availability. In the event of a disaster, it is possible to restore data from one system to the other without losing or missing messages.

7.4 Volumetrics

Volumes of email traffic vary significantly from organization to organization. Analysis of email network traffic and storage requirements is normally undertaken as part of a pre-installation audit.

7.4.1 Storage

The Storage Managers and the Index Managers use permanent storage. Cryoserver's file-based Storage Manager saves messages in compressed files; a compression ratio of approximately 50% can be expected depending on message content distribution together with an overhead of approximately 15-25% for index storage, which depends upon the mix of messages and attachments. Experience has shown that the average storage requirements of a typical user are approximately 0.5GB per year. Cryoserver typically configures its systems with sufficient storage for at least two years, which equates to 1 GB per user. Obviously this varies from one organization and industry sector to another.

Small and medium sized Cryoserver deployments are normally configured using either the C1000 (~50 users) or C2000 (~250 users) appliances. Larger Cryoserver implementations will go for the C4000 (~500 user box) or multiples of, for example a 1500 users will go for 3 x C4000.

For those wanting a customized solution we can utilize and install the Cryoserver software onto HP, IBM, Dell or Sun hardware with securely configured SAN's for data storage. Cryoserver is supported on major SAN platforms including EMC, HDS and HP. In SAN implementations it is possible to use storage replication software such as MirrorView (CLARiiON) or SRDF (Symmetrix) from EMC to mirror data to remote locations without the requirement for a second deployed Cryoserver.

7.4.2 Network Traffic

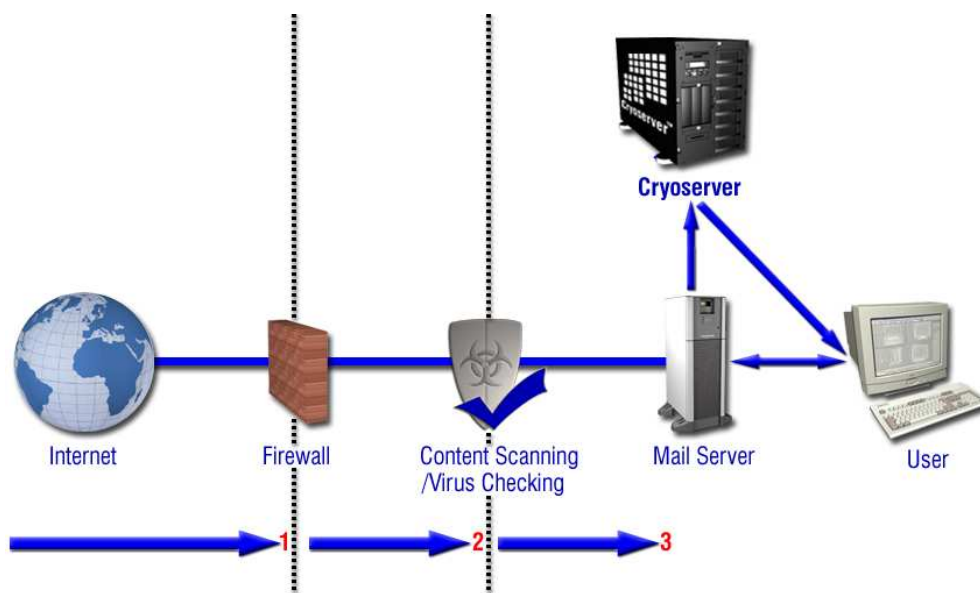
Replicating email messages to Cryoserver will generate some additional network traffic. In general the effect will be to almost double the network traffic due to email messages. However, in most environments the impact on an organization's network infrastructure will be fairly low as the bandwidth generated by email traffic is not normally very high. It is fairly straightforward to estimate the overall effect based on the number and average size of messages sent at peak times.

7.5 Position of Cryoserver in an Existing Network Architecture

The following scenarios depict where the Cryoserver appliance might be placed in an organization's network architecture. Cryoserver is designed to work in concert with existing firewall, virus scanning and content scanning technologies.

Cryoserver was also designed to be as flexible as possible, and so can be integrated with many types of network architecture that are not shown in this document. Large organizations with many mail servers may benefit from having an equal number of Cryoserver devices to reduce network traffic.

7.5.1 Collector Running on the Mail Server



1. Inbound mail from the Internet is allowed onto the network by a firewall.
2. Messages are scanned for viruses and inappropriate content.
3. Messages are delivered to the mail server. Cryoserver takes an audit copy of each message.

This is the recommended configuration because both internal and external messages are captured by Cryoserver. The Cryoserver architecture is flexible and it is possible to configure Cryoserver to meet other requirements.



8. FEATURE & BENEFIT SUMMARY

Regulatory requirements and good business practice can dictate that records be kept for considerable periods of time. There are also new liabilities, and now that the greatest single reason for employee dismissal is email abuse, it makes it even more important to be able to keep full records of emails and speedily find any that need to be relied upon for evidential purposes.

In some highly regulated industries, email must be retained for long periods, particularly in financial, legal and healthcare sectors. In such environments there is now growing awareness that there needs to be a fully auditable forensically compliant copy of individual emails in order to meet legal and regulatory requirements – sometimes in order to show that an organization was *not* in possession of or communicating certain information.

In the event of litigation, records may need to be retained indefinitely. Yet there are costs, time and expense associated with the need for data retention and retrieval. Such systems need to be beyond interference and yet remain easily accessible.



8.1 Table of Features and Benefits

	Features	Benefits
Complete record	<p>Captures a copy of every email</p> <p>Stores a copy of every internal and external email including attachments</p>	<p>Users can quickly retrieve any email that they may have inadvertently deleted (personal disaster recovery) IT department will not waste time trying to restore lost mail Email is captured, even if an employee tries to delete it and hide evidence No reliance upon end-users categorising emails as they decide which ones should be archived. Deters employees from sending non-work related emails. Reduce the amount of emails sent by users, therefore efficiency increases</p>
	<p>Content of emails and email attachments is indexed prior to storage</p>	<p>Allows very fast retrieval at a later date</p>
	<p>No end-user ability to delete emails</p>	<p>Email is legally admissible Fully compliant record. Helps directors and officers avoid accusation of poor record keeping or email shredding. Facilitate cause and effect chain of events to be fully detailed from stored emails Help prove a negative – show what information has <i>not</i> been sent or received</p>
	<p>Intercepts and stores blind carbon copy information</p> <p>Mail servers do not normally record this information, but Cryoserver's record is complete wherever possible</p>	<p>Employees are discouraged from sending confidential information out of the organization using hidden bcc addresses</p>
	<p>Records who was in the distribution list at the time an email was sent / received</p>	<p>Administrator can check who was in a distribution list at a particular time.</p>
Comprehensive audit trail	<p>Privileged and administrator level access to the system is audited using a transcript of their session</p> <p>Transcripts show: log in time, the search criteria, number of results, messages viewed and log out time. These are then stored and sent out to a list of trusted guardians</p>	<p>Ensures employee privacy is protected and access is not abused Protect investigators from being wrongly accused Increases trust in the IT Department</p>
	<p>All privileged and administrative access to the system is recorded and retained securely</p> <p>Summaries of such searches are emailed to nominated trusted individuals.</p>	<p>Ensures that emails are only accessed for legitimate reasons Removes the temptation to snoop on employees Allows officials to monitor each other's activities and enforce good corporate governance</p>
Robust	<p>Mirrored configuration</p>	<p>Complete hardware redundancy allows continued operation even in the face of total loss of a server.</p>



	Features	Benefits
Secure	Emails are stored with a timestamp and “digital fingerprint”	Attempts to tamper with stored data can be detected. Allows proof of an accurate and complete record of messaging data
	Access to the system is only possible through a secure web interface	Only authorised users are able to access other individuals' email. Enforces data access policy. Allows compliance with data protection rules concerning access to personal data
	All traffic passing between the Cryoserver modules is SSL encrypted	Information is protected as far as possible from eavesdropping The front-end is delivered over HTTPS and user passwords are stored in encrypted format
	Level 1- Basic users – Can only see messages they have sent or received Level 2- Privileged users – Can see all messages in the repository. All actions taken are logged for audit to ensure employee privacy is protected Level 3- Administrators – Can create accounts, modify user details and reset passwords. All actions they take are logged for audit. Cannot see any emails. Level 4 - Superusers - Manage overall system configuration and monitor performance and capacity.	All end users can quickly access copies of the emails they have ever sent & received If necessary, and with safeguards, it is possible to find and retrieve any email within the entire repository It is possible to give one department (IT) responsibility for system management (Administrator rights), but give responsibility for email investigations (Privileged rights) to another (HR)
	Emails can be forwarded in real-time to a secondary storage location	Easy to implement long-term offsite storage at a disaster recovery centre All data is secured from accidental or deliberate sabotage by duplication into a trusted secondary physical location Reduce the load on existing mail servers by removing emails
	Email data is compressed before storage	File storage reduced to save on storage space and costs
Flexible	Compatible with current mail servers Microsoft Exchange, Novell GroupWise, Lotus Notes, Sun JES and Teamware are amongst the mail servers supported	Cryoserver will work in an environment with legacy mail systems, capturing copies of emails into a single repository. Investments are protected because Cryoserver allows you the flexibility to change mail server in the future
	Modular system design	Modules within the system may be changed to suit individual requirements, now and into the future.
	Architecture	Repositories can be strategically placed throughout the organization or tactically placed to monitor specific, sensitive business activities either globally or locally



	Features	Benefits
Fast email retrieval	Browser-based user interface	Familiar interface means no end-user training and no need for desktop software installation, cutting down set-up time and costs. Data can be accessed from any convenient location, subject to security constraints
	Simple search on content and metadata	All users can very quickly and cost-effectively track down their own historic emails containing the subject matter that they are interested in.
	Advanced search utilising stemming, sound-alike, and proximity	Retrieve hard-to-find information, whether buried within the body of an email or within an email attachment Allows response to Subject Access Requests within seconds or other disclosure requests in a timely and cost-effective manner
Scalable	More storage and indexing capacity can be added as requirements grow. No impact on the existing storage and indexing regime. Capable of handling the many terabytes of data held by larger organizations	No longer a need for employees to keep insecure and unreliable storage on their local hard-drives (PST files) or to keep within mail quotas The email load on the existing mail system can be reduced, so improving reliability of mail services. Cryoserver can be scaled to suit all sizes of end user organization and service provider.
	Distributed architecture	Network traffic can be reduced by placing storage modules close to individual mail servers
	LDAP / Active Directory support Users can be authenticated against an existing user list.	No need to create or maintain a separate list of email users on Cryoserver – reducing set-up and maintenance costs.