Business Case Analysis:
Email Archiving
for Compliance and Corporate Governance

Ian Andrew Bell,
FrontBridge Technologies, Inc.
ibell@frontbridge.com

Apr 13 / 2004

:: executive summary

Fortune Magazine has referred to email in corporations as “evidence mail”. So ubiquitous is email as the smoking gun in corporate corruption investigations and litigation that companies have begun to provide special training to their employees on standards of acceptable “email behaviour”. However, the cost of settling claims where email serves as crucial evidence pales when compared to the cost of retrieving it from archives when subpoenaed or during discovery processes.

This intricate, painstaking work can take hours per message, can produce reams of data for review and filtering, and since it often occurs at the eleventh hour with extremely short deadlines the result is often huge costs -- often hundreds of thousands of dollars for a single discovery.

The fact of the matter is, nearly all corporations archive email one way or the other for convenience, disaster recovery, or in support of their own litigation or employee disciplinary efforts. The existence of these archives, though, makes them vulnerable to outside investigation, and the existence of such data on linear, offline, removable media such as DLT or DVD-ROM makes retrieval an expensive proposition.

Sarbanes-Oxley places increased importance on the use of email archiving in corporate governance, while rulings such as NASD 3010 and SEC 17a explicitly require messages to be screened and archived. Companies can and have spent millions of dollars in capital expenditure to meet these requirements while others, most notably companies like Porsche, have withdrawn from going public in the United States due to the cost of meeting these requirements.

The opportunity for FrontBridge is to further strengthen its suite of message management services in the Enterprise market. In doing so, and in keeping with the FrontBridge strategy of zero-footprint on the customer premise, a pure managed service represents a credible option both for smaller companies unable to shoulder the burden of complex, online email archiving; as well as for larger companies burdened by the substantial volume and ongoing capital expense of growing and maintaining massive storehouses of email messages.

This opportunity can be realized while maintaining significant margins, being substantially differentiated from competitors in both pricing structure and in the design of the service. By distributing multiple customers across a shared infrastructure in a fully-managed, protected facility, FrontBridge can also enable a high-quality service for small businesses, who traditionally have not had access to such a high degree of robustness.

:: introduction to email archiving

For the past several years, a great deal of attention has been focused on America's corporate email in litigation, regulation, and legislation. As more and more of a company's daily business has come to be transacted over electronic mail, this should be no surprise: According to a Pew Internet study of corporate users conducted in 2002, sixty-three percent (63%) say that email is better than telephone or in-person arrangement of appointments, meetings, and other logistics.

Archiving has taken place for various applications in business for many years. In fact, as a matter of standard practice most businesses archive their email to removable media such as DAT or DLT, and in some cases CD-ROM and DVD-ROM media. This is typically done so that if such messages were lost or the servers that maintained their corporate email experienced a catastrophic failure, there would be some means of recovery and continuity.

But while these archives may prove invaluable in a disaster scenario, their existence makes the company vulnerable to subpoenas and discovery requests as a result of disputes, inquiries, and diligence instigated by outside parties. Regulators and lawmakers have also recently recognized the value of these archives as a goldmine for forensic research into scandal, corruption, and false reporting.

The November 2003 Gartner Group Analysis Report titled "Vendors Respond to New E-Mail Active-Archiving Market Requirements" stated that because of a growing list of regulatory demands, lawsuits, and the implications of the Sarbanes-Oxley Act, that: "Enterprises must begin building an e-mail archive that can be immediately useful for legal discovery and migration of older messages to lower-cost, but searchable, media."

:: email archiving background

As the Enron case dragged on through 2003, the Federal Energy Regulatory Commission (FERC) posted the entirety of Enron's available email archives on its website. This data was accessible to the general public and was famously reviewed by Salon.com among others. What was there served as damning evidence of the abuses of power wielded by Enron executives and the lavishness of their lifestyles at the expense of shareholders, and almost certainly has contributed to the success thus far in prosecuting the company's executives.

More importantly, this validated popular opinion that email messages are a window into the dealmaking, politics, power-brokering, and everyday goings-on of Corporate America. As a byproduct of growing attention to the importance of email and of high-profile debacles such as Enron, then we have seen increasing

attention to the importance of email in substantiating legal disputes instigated by industry regulators, former employees, shareholders, and the financial regulators such as the SEC.

For the Financial Service and Health Care sectors, active policy enforcement and archiving of communication between specific parties within organizations and their outside contacts is mandatory. Sarbanes-Oxley also provides some guidance which punishes management any executive, broadly, who "knowingly alters, destroys, mutilates, conceals, covers up, falsifies or makes a false entry in any record, document . . . with intent to impede." Many companies have interpreted this to construe the deletion of email archives.

With the risk of being "outed" as fraudsters, such as in the Enron case, why do companies absorb the risk of backing up and archiving email? First, because of the mounting pressure by government regulators for fair reporting and auditability of every aspect of a business; Second, and more importantly, because of litigation.

A 2003 survey conducted by the American Management Association revealed that of 1,100 companies polled, 14% had been ordered by a court or regulatory body to produce employee e-mail. The cost of producing subpoenaed or requested documents to support litigation can be enormous. In July of 2003, UBS Warburg paid as much as $300,000 to restore e-mails required for a gender discrimination case. As a result of this precedent-setting decision, defendant corporations are now responsible for bearing the costs of email discovery, assuming plaintiffs can prove that such email-borne information is key to supporting their case.

Given that such information is often widely scattered within organizations -- being sourced from individual workstations, optical media, and/or backup tapes, among other things – the cost of searching, retrieving, and publishing email communications during discovery can cost hundreds of thousands, if not millions, of dollars and require weeks of dedicated effort.

While many companies such as Financial Services and HealthCare Providers have clear mandates to archive their email, as facilitated by regulators, the much broader market will be adjusting to the implications of Sarbanes-Oxley and other corporate responsibility governance.

Smarter companies, driven by their Legal and Human Resources departments, will be searching for solutions which allow them to make cost-effective use of email archives both for internal policy enforcement and for support of litigation.

A November 2003 Gartner Group analysis asserted the following:

> "A new market has emerged, which Gartner has named the e-mail active archiving market, that could be viewed as a subset of the storage

management software archiving market or a hybrid market that spans the storage and content management markets. Products need to include archiving, indexing, hierarchical storage management (HSM) and basic records retention management, as well as special applications for meeting the needs of compliance officers and those tasked with producing records for legal discovery requests. Leading solutions bring together all of these components to address the retention requirements around e-mail while providing the opportunity to more efficiently and cost-effectively manage the impact of growing e-mail data stores."

The marketplace today for email archiving consists primarily of software vendors, and much less often service providers, teaming with WORM-compliant storage vendors such as EMC's Centera system or actual WORM storage managed by companies such as Iron Mountain.

The vast majority of effort in this market so far is by software vendors selling solutions to Large Enterprise, mated to specific email MTAs, as this Gartner Dataquest table shows:

Table 1
E-Mail Active-Archiving Vendors Shipping Product in Mid-2003

Vendor	Product	Exchange	Domino	Others	General
AXS-One	AXS-One Email and IM	X	X	SMTP[1]	2003
Connected	MailStore[2]	X	-	-	2002
Educom TS	EAS	X	-	-	2000
EMC Legato	EmailXtender	X	X	Sendmail	2000
Entelagent	SAMS Online	-	-	SMTP	1998
IBM	CommonStore	X	X	-	2000
iLumin	Assentor	X	X	SMTP	1998
IMR	MailStore	X	-	-	2002
IXOS	ECONserver	X	X	-	2000
KVS	Enterprise Vault	X	-	-	1999
Nexic	Enterprise Discovery	-	-	GroupWise	2003
Sun	Infinite Mailbox	-	X	-	2003
TJ Group	CAI Suite	X	X	-	2000
Tower Technology	IDM eMail Archive Option	X	X	SMTP	2003
Tumbleweed	Message Monitor	-	-	SMTP	1998

[1]Captures all e-mail at the SMTP Internet mail relay or gateway, e-mail application independent
[2]Connected will rename its product ArchiveStore in 2004 with the next version release.
Source: Gartner Dataquest (November 2003)

:: market opportunity

While archiving for compliance touches a limited subsection of the Financial Services sector, healthcare providers, and regulated utilities there is a compelling drive toward archiving for corporate governance driven by the increasing demands for outside scrutiny, and the massive costs associated with retrieving data from existing archiving practices under pressure from Discovery or other outside scrutiny.

While the solutions predominant in the marketplace today are well within reach of Fortune 1000 companies, there are precious few (arguably, no) choices for companies in Small-to-Medium Business. Additionally, the cost of deployment represents a significant capital investment even for larger organizations, and many would clearly be open to a service model which has costs that scale directly in line with usage.

FrontBridge cannot immediately enter the market for Compliance Archiving, since such a product would also require realtime message management for Compliance. While this is on the roadmap, it has yet to be thoroughly considered for development and subsequent release.

In the interim, it was decided to pursue various “feet wet” strategies in developing a low-cost managed archiving service for email for corporate governance.

:: product design

At the highest level, our initial service will be a fully-hosted email archiving and retrieval service, allowing companies to journal from their MTA into the high-reliability FrontBridge network, where messages will be archived securely onto non-rewriteable but highly accessible media and can be retrieved via a web-based interface by duly authorized personnel, using powerful search capabilities, easily and cheaply.

As such, it is desirable for the service to operate as an application layer on the FrontBridge mainline network both to enable efficient cacheing, queueing, routing, and high-performance apparent to the sending MTA. This will be a key demonstrable advantage for FrontBridge over any other solution in the marketplace.

:: partnership opportunities

It would be difficult, given current human resource constraints, for FrontBridge to build a credible technology for entry into the marketplace. Therefore, the build vs. buy decision was made very early in the analysis for this service. Almost

immediately we began discussing our interest in the product with potential partners. There are two key elements to building a hosted archiving solution that are critical: Archiving Software, and the Storage Medium.

Archiving Software

The first key element is software which will enable journaling from the MTA, indexing, search and retrieval, and management on messages which are stored. Various partnerships were explored, including KVS, which ranked as number one in Gartner's recent Magic Quadrant. Gartner's Magic Quadrant rated players in this sector and yielded the following results:

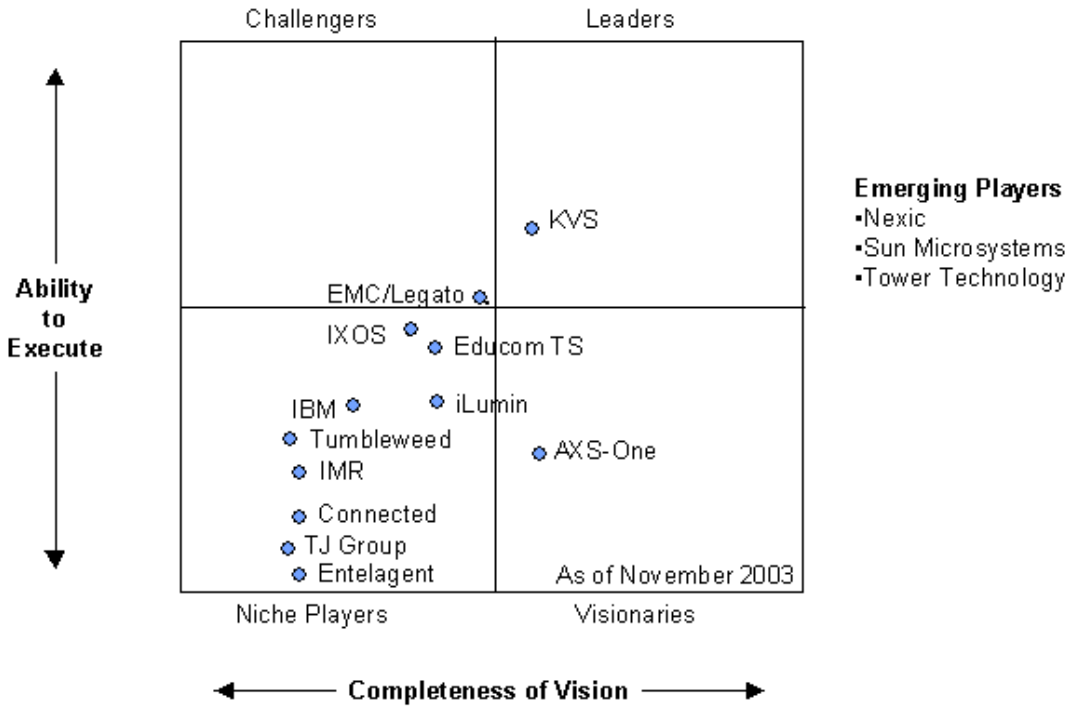

Notably, many key market players were not included in this research without adequate explanation. In the case of FrontBridge, which has requirements peculiar to the Managed Services Reseller market, several key evaluation criteria were applied to vendors, both those mentioned in this report and sourced through other research:

Vendor:	KVS	iLumin	Zantaz	ZipLip	Legato
Multi-Tenancy:	No	No	No	Yes	No
SMTP:	No	Yes	Yes	Yes	No
Notes/Domino:	No	Yes	Yes	Yes	Yes
Exchange:	Yes	Yes	Yes	Yes	Yes
OS Platform:	MSFT	MSFT	Linux	Linux	MSFT
Scalability:	Medium	Low	Medium	High	Low

Key requirements for FrontBridge were support for multiple MTA platforms, multi-tenancy, and ideally the partner's platform would be based on technologies similar to those utilized today by FrontBridge.

While Zantaz has a very significant, credible product offering that would fit well into the FrontBridge Architecture, they were deemed to be competitors. Similarly, discussions with iLumin were inhibited by an announced partnership with Tumbleweed and their inability to support multi-tenancy.

An intangible here is the use of the partner as a lead-generation tool. During discussions with iLumin it was contemplated that iLumin could act as a referrer of new business to FrontBridge. Unfortunately ZipLip and others could represent no such similar advantage.

In detailed discussions it has been ascertained that at present a partnership with **ZipLip** represents the best chance for success in going forward. Trials will be soon underway to demonstrate the software contiguous with the systems integration of the elements.

Storage Medium

The second is the storage medium. In order to meet strict legal definitions of non-rewriteable media and later to comply with regulatory requirements, a SAN-based WORM was deemed necessary. While this technology is relatively simple to build using open-source building blocks and commodity hardware, it was decided that there was significant advantage in building the service's credibility by leveraging the market's understanding of an existing, off-the-shelf product which had already been accepted by the SEC for compliance.

Easily the most well-known product in this small but rapidly developing marketplace is the EMC Centera. As the diagram below portrays, the Centera and products similar to it are Storage Area Networks, which in this case is comprised of commodity, Intel-based PC servers running Linux.

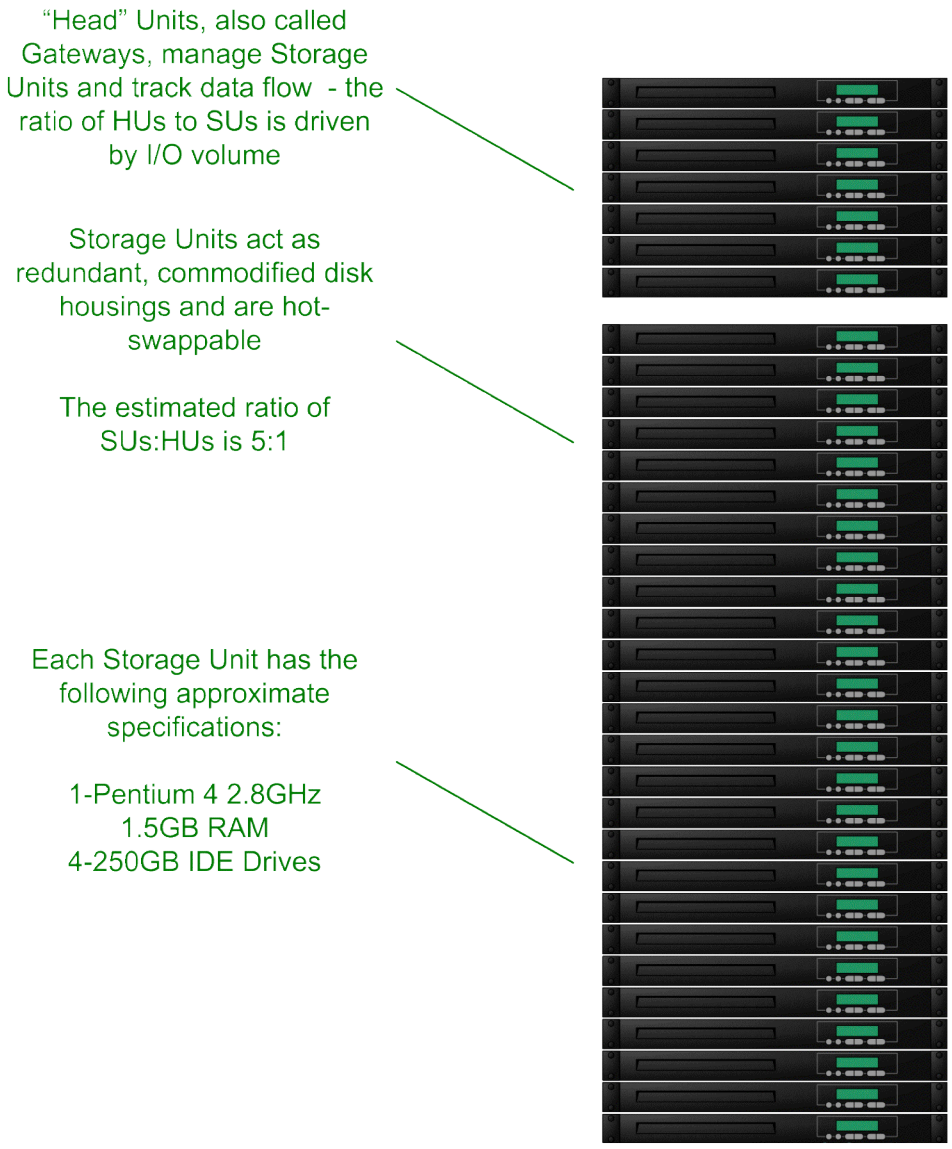

An 8-Node EMC Centera system consisting of 2 HUs and 6 SUs represents 3TB of usable disk space costs approximately $105,000.00 USD. The Centera is a proprietary solution consisting of software developed by EMC and hardware sourced, integrated, approved and resold by EMC. Centera has been deployed numerously for compliance in major Enterprises and, based on FrontBridge's existing relationship with EMC the company has been very aggressive on pricing.

The Centera communicates with the Archiving Software via a proprietary API. As yet, ZipLip has yet to integrate to the Centera API, however EMC claims that this effort is trivial and that sentiment has been echoed by ZipLip technical staff.

A second potential partner was sourced during the investigation into WORM-on-disk (also known as Content Addressed Storage) solutions, Permabit. Permabit's Permeon Compliance Vault is a software-only solution that would

allow FrontBridge to repurpose its own hardware, or hardware purchased under existing supplier relationships, for storage. As an advantage this reduces the capital risk of deploying the service, allows for common maintenance spares between FrontBridge's existing infrastructure and the Content Addressed Storage system. Another significantly appealing attribute of Permabit is its flexibility – Permeon can be configured to meet SEC-17a compliance or can be used to drive a more vanilla storage application. This means that, using the same software licenses, FrontBridge could adapt a CAS network driven by Permeon to meet its changing product goals.

Permabit recommends the following configurations, the latter of which represent systems already used by FrontBridge in other applications:

```
==============================================
Area Electronics (white box)
----------------------------------------------
Portal          NAS-908          1 x 200GB    1U
Storage         NAS-908          4 x 300GB    1U
==============================================
Dell
----------------------------------------------
Portal          Poweredge650     1 x 80GB     1U
Storage         PowerEdge2650    5 x 146BG    2U
==============================================
Hewlett Packard
----------------------------------------------
Portal          DL320-G2         1 x 80GB     1U
Storage         DL380            6 x 146GB    2U
==============================================
IBM
----------------------------------------------
Portal          xSeries 305      1 x 120GB    1U
Storage         xSeries 345      6 x 146GB    2U
==============================================
```

A comparative analysis reveals significant advantages with Permeon, they have yet to officially test acceptance by SEC 17a. According to their sales personnel they are presently awaiting the negative response of the SEC with a customer who has submitted Permeon as a cornerstone of their compliance plan, which process should be completed in May.

The following chart outlines the comparatives between EMC and Permabit:

	EMC Centera	Permabit Permeon
Entry Cost	~$105K	~$75K
Credibility	HIGH	LOW
Flexibility	LOW	HIGH
Infrastructure Fit	LOW	HIGH

Interestingly, while the Permabit solution using IBM equipment similar to that deployed by FrontBridge is somewhat cheaper, it uses approximately double the

rack space of the Centera offering. The increased cost of IBM hardware in this scenario is frustrating to the “cheaper” facet of the value proposition.

Additionally, since Permabit is exposed to the Archiving Software as an NFS-mountable volume, special provisions must be taken to ensure true multi-tenancy. Permabit has promised a PERMEON software revision in May which supports this. Since EMC’s Centera is API-driven, and not a mountable volume, its integrity and security is on a per-bit basis and therefore tenancy context is not applicable.

:: service design

Leveraging the FrontBridge Architecture is key to the success of this service, and as always, will represent a key advantage versus competitors who will presumably follow. The FB architecture also provides a key cacheing, queueing, and load distribution mechanism which naturally lends this service much of its reliability, performance, and redundancy.

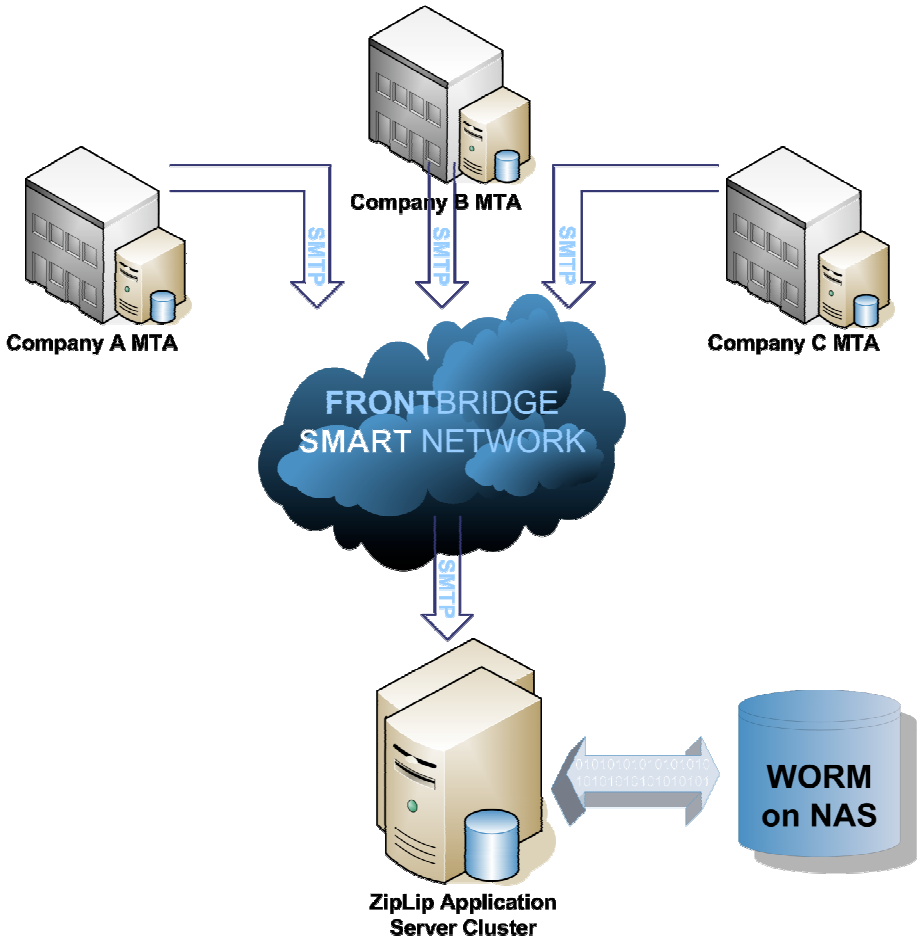

By journaling all email directly from the Mail Transfer Agent (MTA) at the company site, FrontBridge ensures that it is able to capture both the company's internal messaging and mail flowing into and out from the corporate domain. By leveraging the robustness and automatic failure handling of the SMTP protocol, journaled entries will be addressed to an arbitrary destination account along the following schema: archive@company.archive.frontbridge.com. Similarly, when compliance or other staff wish to query the archive they will enter the web portal at http://company.archive.frontbridge.com. This simplified schema will allow FrontBridge to virtualize the hosting service, allowing it to be moved to different logical or physical ZipLip application servers within the Application-specific network.

:: financial analysis

Estimates of traffic based on Pew Internet's conservative study of bandwidth consumed by email reveal that the average business email user generates 1.75MB of message traffic daily, comprising incoming and outgoing email. Based on this statistic, 10,000 users would consume approximately 6 terabytes (TB) of data per year in email storage.

With a relatively minor initial investment in infrastructure, FrontBridge could deploy an 8-node EMC Centera system, capable of storing just under 6TB, for approximately $105,000.00 USD. ZipLip pricing has yet to be negotiated but is based on a total cost of $361,533.00 for 20,000 seats -- inclusive of all software with 3 years of maintenance included.

The simplest pricing strategy in the industry, as gleaned from Iron Mountain, is to charge a setup fee per seat in addition to storage.costs on a monthly basis. This makes it easiest for customers to calculate their expected costs for the service and to make intelligent choices about the length of retention, and allows FrontBridge to balance its risk and continually make smart capital investments directly in line with growth of the service (and thus its data storage appetite).

Proposed Pricing:

Seats	Cost	Margin 3Y	Margin 2Y	Margin 1Y
0-500 seats	$25 per seat signup $11 per GB per month storage	**84.5%**	**73.66%**	**41.89%**
501-5000 seats	$20 per seat signup $10 per GB per month storage	**82.8%**	**70.60%**	**33.46%**
5001-10000 seats	$15 per seat signup $9 per GB per month storage	**80.7%**	**66.74%**	**22.18%**

* fixed costs distributed across 20,000 seats

…this represents a substantial margin opportunity for FrontBridge in marking up storage costs and creates elasticity in competitive pricing situations. As the table above displays, the key to the success of the service rests on one of two scenarios:

- Customer retention for at least 24 months, or
- Customers migrating to the service with pre-existing data

A mailbox which has been archived by FrontBridge for 36 months will pay a monthly fee on their third anniversary of approximately $21 to host 1.8GB of email. This pricing is actually competitive (at a slight premium) over email hosting services which archive mail on non-Centera storage media.

For cost recovery, the first 10,000 Seat-Years (12 months of service for 10,000 seats) will garner approximately $250,000.00 in signup revenue, and $399,902.34 in revenue for storage, for a cumulative take of $649902.34.

In contrast, the maximum CapEx for launching the service (assuming no terms have been negotiated with ZipLip) will reach as high as $494989.00. Analyzed from this perspective, the service would achieve a margin of 24%, however this figure is deceptively low. Total payback for the service, however, is relatively easily achieved within the first year of operations, well within the first 10,000 seats.

It is very likely possible to negotiate a graduated licensing schema with ZipLip that more appropriately balances risk/reward. It is possible that the service could be fully launched for less than $200,000.00 USD.

(See accompanying spreadsheet for formulae and calculations).

:: competitive options

The vast majority of options available to customers in this space are of course designed for Large Enterprise, requiring significant capital investment. Some solutions companies, such as Iron Mountain, Zantaz and Tumbleweed, will manage dedicated infrastructure on behalf of the customer. While some economies of scale are realized by sharing the management resources and expertise of Zantaz personnel there is little net cost benefit to hosting vs. building and operating an infrastructure. Like in email protective services, companies will make a choice between in-house deployments and a managed service and such managed services must have significant advantages from both a cost and an operational perspective.

Provider	Cost
Iron Mountain	$30 per seat one time $12 per GB per month storage
ArchivedEmail.com	$45/mo. (incl 1GB), $12.50/GB transfer $12.50/GB per month for storage
Zantaz	$0.18/MB one-time, included storage duration unknown

Other providers in this space include: Epsiia, Seccas, and SectorInc.

At Zantaz's $0.18/MB price, the first-year bill for a 200 seat company will be $146,952.00, versus a cost to a FrontBridge customer of $15,264.00. Clearly our product as conceived makes the service accessible to a smaller market.

:: swot analysis

A core strength of this product offering is that it enables access to a robustness and level of compliance to a customer set that it was previously unavailable to. Small- To Medium- Enterprise is subject to many of the same motivations and requirements as Large Business, but solutions (because of infrastructure cost) have been entirely targeted at Fortune 1000 companies. Arguably, the only way small- and medium-sized companies would have access to such a service is with a hosted model, capturing economies of scale realized by distributing infrastructure and management costs across multiple distinct organizations.

With this strength comes a price: risk. Few companies have attempted to build a shared-infrastructure archive solution such as this – Service Providers have previously attempted to offer managed services by building up and maintaining dedicated software and equipment on a per-customer basis. Ours will be one of the first attempts to balance multiple customers on a single shared infrastructure – however this is a very familiar story to that which FrontBridge encountered when entering into the message management space.

Opportunities for this service abound – from moving from the initial launch phase as a message archiving service for corporate governance to full 17a compliance, to offering a full disaster recovery and business continuity service based on a more active email archive. Most importantly, this divergence from pre-delivery message management to post-delivery message management heralds a more comprehensive range of solutions, conditioning the customer to understand FrontBridge's role as a general Enterprise Messaging Services company. This service will benefit from a first-mover advantage and, like existing anti-spam and anti-virus services, is a cure for pain which businesses feel less so than it is a forward-looking new business application.

The major threats to this service are fairly common. There is in our sector considerable risk of downward price pressure as has been experienced with email protective services. Additionally, there is some liability risk in storing customer data in a secure data centre, should that single site be subjected to fire, earthquake, or other significant event which destroys that data. On the other hand, the latter risks are fairly standard and are covered under Force Majeure exemptions in standard SLAs. A threat which may inhibit the product's growth is rejection by the SEC in a situation where a customer subjects this to the 17a compliance submission process. This would essentially force the company "back to the drawing board" for enhancing this product to meet 17a. Only through consultation with subject matter experts (one such credible expert has been identified) can FrontBridge mitigate this possibility.

:: conclusions

FrontBridge should proceed to immediately develop and launch this service. The impact from a development and operational standpoint is negligible, entry costs can be mitigated to reduce risk, and there are not many alternatives for the target market. The early value proposition of Corporate Governance should be pursued until it is sufficiently established that FrontBridge's Email Archive service can meet SEC 17a requirements (this should be done with a friendly customer).

Systems integration with the FrontBridge SMART architecture is a key advantage which this service will represent over others in the market – and that functionality is essentially made available for free.

According to calculations, the bandwidth impact of even a large number of users will be negligible both on the customer end and for the FrontBridge architecture is minimal. Customers with over 3,000 seats may opt to journal via a dedicated DSL connection, however with sufficient backbone access their journaling traffic will likewise be a small percentage of overall traffic.

SSL/TLS will also be critical, for customers who do not wish for their journaling stream to be monitored by third parties.

:: attached documents for review

- **Gartner Research Document**
- **Compliance Archiving Discussion Paper**
- **FB Archive Business Case Analysis Worksheet**