The Importance of Email Continuity

An Osterman Research White Paper

Published December 2010

SPONSORED BY

Osterman Research, Inc. • P.O. Box 1058 • Black Diamond, Washington 98010-1058
Tel: +1 253 630 5839 • Fax: +1 253 458 0934 • info@ostermanresearch.com
www.ostermanresearch.com • Twitter: @mosterman

Executive Summary

OVERVIEW

Email is the single most important tool in helping employees to get work done in organizations of all sizes. Most email users have become so dependent upon their corporate email system that they check email almost constantly while at work, most check it at home on weekdays and weekends, they check it while traveling and also while on vacation. The average email user spends nearly 30% of his or her day doing something in their email client or browser-based equivalent. And, despite the conventional wisdom, email is being used more over time – not less – as it becomes the focal point for integrating various types of communication and business processes.

However, email is not solely about communications – just as important as access to email for *communications* is its role as a *repository* of critical business information – a role that finds email systems storing a greater proportion of the content that users and their employers need for both the short- and long term. For example, a May 2010 survey by Osterman Research found that 55% of corporate email users who employ email up to two hours each day consider themselves "pack rats" or nearly so in the context of storing information in the email system; 72% of those who use email more than two hours each day are such storers of information in email.

KEY TAKEAWAYS

There are two key takeaways in this white paper:

- Because access to email and the information in email systems is so critical, it must remain continuously available. Outages of even a few minutes in length can create numerous problems for individual email users, as well as the business processes that rely on email as their transport infrastructure.

- Email systems that go down cost organizations in a variety of ways, including the loss of employee productivity, extra IT time investments in resolving email outages, lost business opportunities, and a variety of less tangible consequences.

As a result, every organization should implement an email continuity solution to ensure that email remains operational as close to 24x7 as possible. A failure to do so can have serious repurcussions that will be felt in both the short and long term.

ABOUT THIS WHITE PAPER

This white paper discusses the importance of deploying an email continuity solution and provides advice on the key issues that should be considered when doing so. It also provides a brief overview of Google, the sponsor of this white paper and a leading vendor of email continuity solutions.

The Importance of Email

EMAIL IS WHERE USERS SPEND MUCH OF THEIR TIME

Email has evolved from a tool focused primarily on communications to something of an information portal for the typical user. For example, most email clients or browser-based equivalents are used to:

- Send and receive email
- Send and receive word processing documents, presentations and spreadsheets
- Create, respond to and be reminded of appointments
- Manage tasks
- Manage contacts
- Manage real-time communications
- Store documents of various types
- Take notes

Further, email is used increasingly as a portal for social networking interactions, a sort of clearinghouse for various social media feeds, not to mention the integration of real-time communications into email clients that is replacing standalone instant messaging clients.

However, email systems are also the communications backbone for various types of automated information delivery systems, including those focused on supply chain management, customer relationship management, transaction processing and other critical business processes.

In short, email has become the integration point for users' communication experiences, the primary method for sending files to others inside and outside the organization, and the most commonly used collaboration platform. Email has also become a critical backbone for timely delivery of information sent between systems or from systems to users.

HOW MUCH IS EMAIL USED?

It is no surprise, then, that email is the single most used application for the typical corporate user. In a survey conducted by Osterman Research during November and December 2010, we found that the typical user spends 133 minutes per day working in email, or about 28% of an eight-hour day. This is dramatically more time than is spent using the telephone (61 minutes), using real-time communications tools (28 minutes) or using social media (11 minutes). Further, the typical user sends a mean of 43 emails on a typical day and receives 123.

It is also important to note that email is used widely on smartphones and other mobile devices, further increasing its use by the typical user.

EMAIL IS BECOMING MORE CRITICAL OVER TIME

Despite the proliferation of new communication and collaboration platforms – and despite pundits' forecasts to the contrary – email use is increasing over time, as it becomes the focal point for integrating various types of work processes. It is also important to note that email is often the preferred method of communication or file transport even when purpose-built tools are available. For example, despite the availability of FTP and other tools that often are a more efficient method for sending large files, users still typically rely on email as the primary method for sending contracts, proposals, orders and other business content. Users often will use email as a real-time communications tool when in a meeting, often sending a document and waiting for the recipient to receive it before continuing the discussion.

The bottom line is that email is critical to the way that people work and it is becoming more so over time, both in terms of the critical nature of the work processes that rely on email and in the amount of time that people spend using it.

What Happens When Email is Not Available?

CURRENT LEVELS OF DOWNTIME

Overall, corporate email systems are *somewhat* reliable: an Osterman Research survey conducted in August 2010 found that email systems in mid-sized and large organizations in North America experience a mean of 37 minutes of unplanned downtime during a typical month in addition to a mean of 84 minutes of planned, maintenance-focused downtime each month. This translates to system uptime of 99.72%.

However, such *seeming* reliability masks the fact that 26% of organizations experience more than 30 minutes of unplanned downtime each month – in fact, 12% of organizations experience two or more hours of unplanned downtime each month. The distribution of email downtime in mid-sized and large organizations is shown in the following figure.

Amount of Email Downtime Experienced During a Typical Month

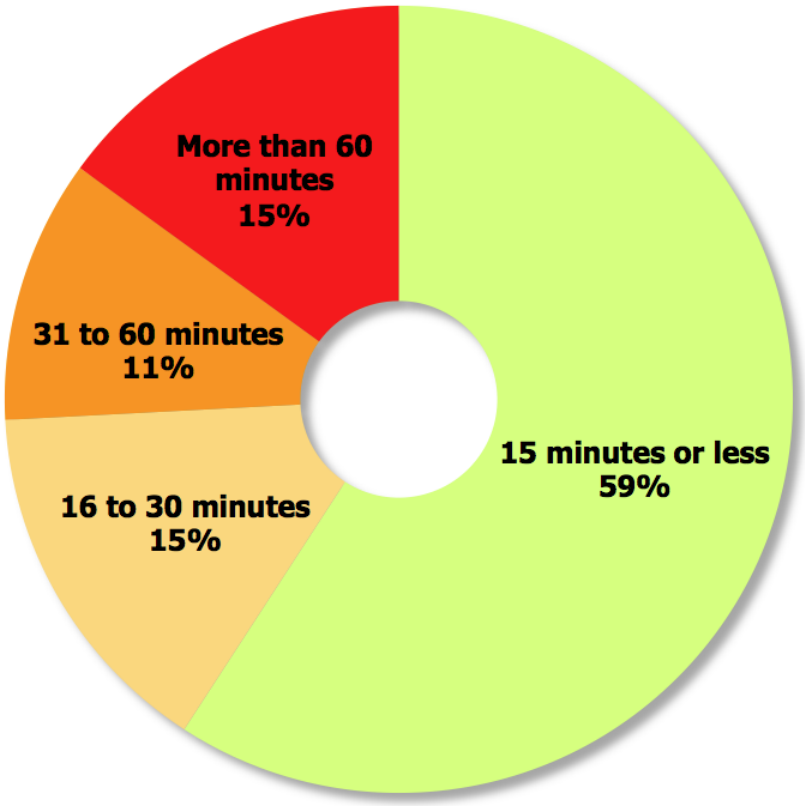

COSTS ASSOCIATED WITH DOWNTIME

Email downtime carries with it a number of consequences that can vary widely in their short- and long-term impacts on an organization. For example:

- If we assume that email users are just 28% less productive during an email downtime incident (the proportion of their day spent using email), a very conservative figure of 30 minutes of unplanned downtime each month for users whose fully burdened salary is $65,000 per year will result in an annual productivity cost of downtime totaling $52.50 per user per year – an organization of 500 users will, therefore, suffer productivity loss of $26,250 each year.

- However, productivity costs can be substantially greater if email is unavailable for extended periods, as can occur during power outages or storms. This might result in productivity loss of 100% if employees go home for the day, resulting in significantly greater productivity losses than in the example above.

- Of particular note in the context of downtime are mobile users who typically are more sensitive to email interruptions. This sensitivity to downtime – and the economic consequences that accompany it – are driven by the fact that mobile users often have no viable alternative means of communication if they cannot

communicate via email, they are often more pressed for time while traveling, and they have shorter windows in which to communicate (e.g., while on a layover at an airport).

- During outages of the corporate email system, users will look for alternate means to send time-sensitive or other critical content, including personal Webmail systems. Use of these personal systems means that corporate data is sent without first being processed by the archiving, content filtering, security, encryption or other corporate systems that might be in place, putting the organization at risk on a number of levels. This cost is particularly difficult to quantify, since it can result in consequences as diverse as leakage of sensitive corporate data all the way to charges of spoliation of evidence in a lawsuit.

- There are a number of other difficult-to-quantify consequences from email downtime, as well, including loss of corporate reputation when prospects, customers and business partners receive a bounceback when sending email – a prospective customer who sends an email only to have it bounce back might never follow up later. An email outage might prevent a timely response to a proposal, order or answer to a client inquiry and could lead to a significant loss of revenue in the long run. Unplanned downtime incidents require IT staff to stop doing other work and instead focus on the email emergency at hand, thereby delaying other work.

NATIVE EMAIL CONTINUITY CAPABILITIES OFTEN ARE NOT ADEQUATE

For many organizations, resolving unplanned downtime incidents is a time-consuming process. Osterman Research has found that 25% of organizations require anywhere from 1.1 to four IT person-hours to resolve the typical unplanned downtime incident, while 17% require more than eight IT person-hours to resolve such an incident.

If we assume an average of just 2.0 IT person-hours to resolve a single, unplanned downtime incident in email, and that there is an average of only one such incident per month, that means that IT will spend 24 person-hours each year dealing with downtime incidents. However, this represents something of a best-case scenario, since many IT departments spend much more time than this resolving email downtime.

A key contributor to the length of unplanned downtime incidents is that native email continuity capabilities are often inadequate. The length of time to detect, diagnose and remediate email downtime is often excessive, leading to email downtime incidents that are longer than they should be. Exacerbating the problem is that quite often end users are the first to notice an email outage, not IT. This adds to the length of time required to resolve downtime incidents, since IT diagnosis and remediation efforts often do not commence until at least 10 minutes after the email system has gone down. If we assume that an email system goes down once each month, this represents at least two hours of downtime each year attributable only to the fact that users are the "canaries in the coal mine" of email downtime detection.

THE IMPORTANCE OF SLA, RTO AND RPO

In understanding the requirements for email continuity, it is necessary to determine exactly "how much" email continuity is necessary to meet the needs of individual users

and the organization overall. Toward this end, there are three critical issues on which to focus in the context of email continuity:

- **Service Level Agreements (SLA)**
  A Service Level Agreement is simply a formally defined level of service that will be provided by an IT department or a service provider. For example, an email SLA of 99.9% means that an email system will be down no more than 43 minutes per month.

- **Recovery Point Objectives (RPO)**
  It is important to determine just how much email data loss is acceptable following an outage. An organization may determine that it can afford to lose several tens of emails per user following an outage, and so can afford to establish an RPO that follows the outage by a substantial length of time. However, most organizations will likely find it unacceptable to lose a substantial amount of email, and so will want to establish an RPO that extends only to shortly after the outage begins.

- **Recovery Time Objectives (RTO)**
  Related to RPOs and SLAs is the RTO, which involves determing how much time between the commencement of an outage and recovery is acceptable.

While RPO and RTO often are more focused on backup technologies than email continuity, many consider them to be important considerations for a recovery capability, as well.

Key in determining RPO and RTO for a particular organization is determining the benefits of reducing or eliminating data loss, as well as determining the benefits of speeding up the restoration of email service following an outage. The SLA, RPO and RTO will be based on a number of factors, including:

- **Specific employee productivity requirements**
  Organizations whose employees rely heavily on email to do their work will normally be fairly intolerant of downtime, and so SLAs, RPOs and RTOs must be carefully defined.

- **The processes and systems that rely on email as a transport mechanism**
  If supply chain workflows, call center transactions, revenue-generating operations and other critical processes rely on email to transport content, the SLA, RPO and RTO requirements will be very different than for an organization in which email is used only for person-to-person communication.

- **Legal and regulatory requirements**
  Email outages can create gaps in the content archiving process, causing major problems for organizations that are heavily regulated, such as those in the financial services, energy or government space. Minimizing email downtime in these cases may have less to do with bringing communications back online than ensuring that critical business records in email are retained appropriately.

The following figure provides an overview of the relationship between SLAs, RPOs and RTOs.

Relationship Between SLA, RPO and RTO Based on Email Requirements

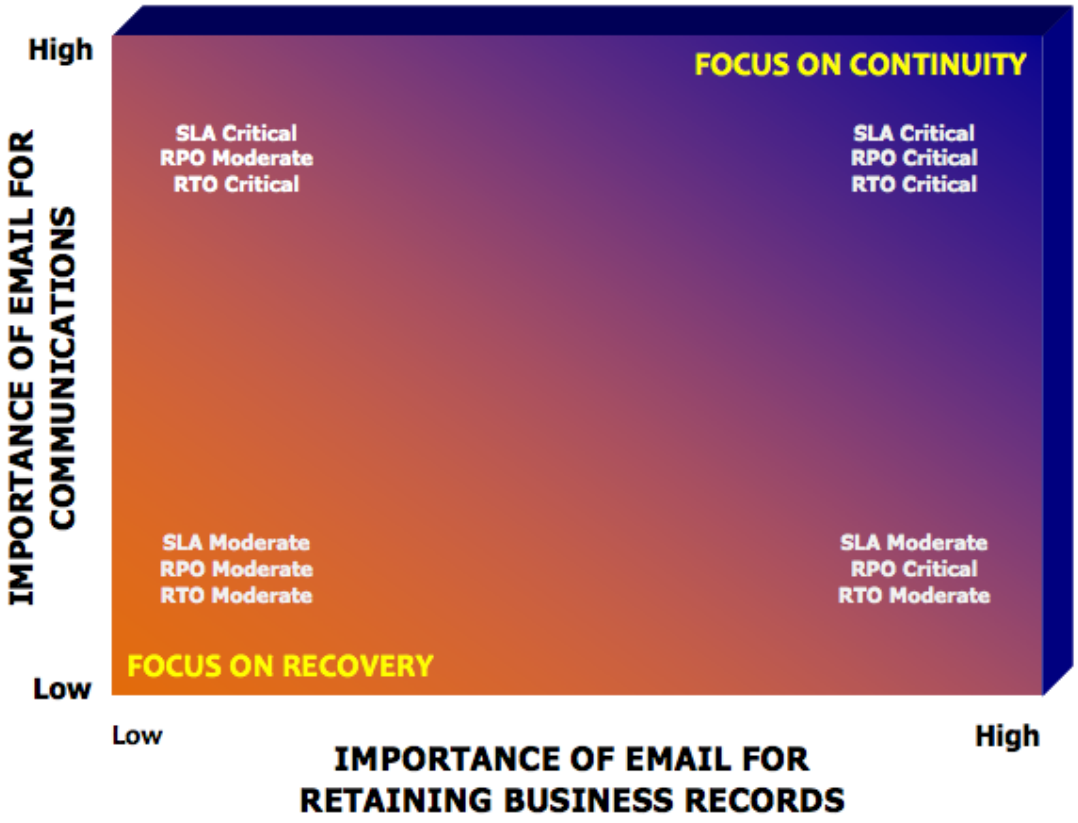

What Should Be Your Next Steps?

UNDERSTAND THE TRUE CONSEQUENCES OF EMAIL DOWNTIME

The initial step in creating an email continuity strategy is to understand the risks that an organization faces from when email suffers an outage. These risks might include:

- **Lost employee productivity**
  Employee productivity is a "soft" cost – because senior management does not write a check to cover this cost, they may be tempted to disregard it. However, it is a real cost of email outages and should be considered when determining the ROI for an email continuity capability. As noted earlier, the cost of lost productivity, even for a relatively small firm, can be significant.

- **Natural disasters**
  The likelihood of an email outage caused by a natural disaster occurring in the areas in which you do business and in which your email servers are located (assuming

your email capabilities are not provided by a cloud-based provider) is very high over the long term. Of course, an organization with operations in earthquake-, hurricane- or tornado-prone areas will suffer more email outages than those in more benign environments, but eventually every organization will face an email outage caused by natural forces.

- **The regulatory environment**
  Regulations that directly impact an organization can also play a key role in defining email continuity requirements. For example, a firm that is required to preserve communications with clients in a serialized fashion, such as broker-dealers, faces a much more serious risk from an email outage than a firm that uses email for basic communication and that does not store business records in email.

- **Email downtime varies by platform**
  The email platform that an organization has in place can also have an impact on the amount of downtime they experience. Osterman Research's annual survey focused on messaging servers highlights significant differences in the average amount of downtime experienced by leading email servers.

DEFINE AN ACCEPTABLE SLA, RTO AND RPO

Every organization must establish an acceptable SLA, RPO and RTO when defining their email continuity plans. For example:

- An uptime of 99.9% will result in downtime of 43 minutes per month or eight hours 46 minutes of downtime annually. It is important to evaluate vendors' claims of reliability with careful scrutiny – a claim of 99.99% uptime means that the system experiences only about four minutes of downtime each month, a level of performance that is particularly difficult to achieve in many email server environments.

- An RTO, while similar in concept to the SLA, defines the objective for the maximum length of time that should transpire between the commencement of an outage and its resolution. Unlike the SLA, the RTO does not define the minimum amount of uptime during a given period.

- An RPO defines the maximum acceptable time between an outage and the restoration of data from the downtime incident. It does not necessarily when the email system has been returned to normal operation. For example, if the RPO is 10 minutes, then the organization has defined that an outage can create no more than 10 minutes of data loss. This does not mean that the system will be operational within 10 minutes – it could be down much longer than that – but that no more than 10 minutes of email data will have been lost as a result of the outage.

DETERMINE THE ROI OF EMAIL DOWNTIME REMEDIATION

Most organizations should probably conduct an ROI analysis for their email continuity requirements. For example, assuming an email continuity system is priced at $20,000 initially, plus an additional 20% per year for maintenance and one hour of IT labor per week to manage it (and 5% salary growth per year), the three-year cost of ownership

for the system will be $34,305. If the system will save a 1,000-seat organization $52,500 per year in employee productivity losses and will enable the organization to earn $25,000 per year in additional revenue, the ROI will be 601%, calculated as follows:

(Benefits – Cost) ÷ Benefits = ROI

($240,506 in benefits over three years - $34,305 in system cost) ÷$34,305 = 601%

The same type of ROI analysis can and should be performed for cloud-based services. ROI analysis may actually be slightly easier for cloud services, since their exact cost is defined for the length of the contract period.

It may not be necessary to conduct a thorough ROI analysis in every situation. For example, email is absolutely critical to most users and organizations, and so qualitatively defining the benefits of maximizing email uptime may be sufficient for some decision makers to justify deployment of an email continuity solution.

IMPLEMENT AN EMAIL CONTINUITY SOLUTION TO ENABLE RAPID FAILOVER

Finally, it is important to deploy a robust email continuity solution that best meets the needs of the organization. For on-premise solutions, decision makers will need to determine their SLA, RPO and RTO and choose an appropriate solution. The options include:

- **Basic backup systems**
  Backup systems can provide a reasonable RPO or RTO, particularly if they are disk-based, but will often be lacking in terms of SLA.

- **Replication/failover**
  These systems generally offer better SLAs, RPOs and RTOs. Some may opt for cloud-based solutions to provide a backup email capability using the corporate domain. These solutions generally are operational very quickly after an outage in the primary email system.

- **True email continuity solutions**
  These solutions are designed to provide continuous availability with no disruption to end-users. This is clearly the best option because it imposes little or no interruption in the use of email. It is also important that an email continuity solution provide continuous availability to all of the systems that form the messaging infrastructure, including email servers, mobile device servers, security servers, etc.

The delivery model must also be considered as part of the solution planning as shown in the following table that presents the pros and cons of each approach.

Various Delivery Models for Disaster Recovery and Business Continuity

Delivery Model	Pros and Cons
Software	• Can be the least expensive approach, particularly if existing servers can be repurposed • Can be more expensive to manage because of the use of in-house IT staff and/or if remote data centers must be maintained • Widest variety of solutions available • Can be more expensive than cloud-based solutions if the architecture chosen to support a large number of geographically distributed sites each require their own email continuity infrastructure
Appliance	• Can be less expensive to deploy than software-based solutions when IT evaluation and deployment costs are added in • Can be faster and easier to deploy than software-based solutions • Can be more expensive than cloud-based solutions if the architecture chosen to support a large number of geographically distributed sites each require their own email continuity appliance
Cloud-based service	• Normally the most rapid and easiest solution to deploy • Easiest option for internal IT staff because a third-party manages the solution • Can be less expensive to deploy and manage than on-premise solutions, particularly for smaller organizations • Reduces on-premise power requirements • May be more expensive than on-premise solutions for large organizations

Sponsor of This Report

Google, Inc.
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
+1 650 253 0000
www.google.com

Google Message Continuity, powered by Postini, is a cloud-based disaster recovery solution for organizations running Microsoft Exchange email servers. By providing Gmail as an alternate, synchronized email system, Google Message Continuity delivers an RPO design target of zero and an RTO target of instant failover. With the reliability and security of Google's services extended to your on-premise environment, Google Message Continuity allows you to:

- Develop a complete email continuity and disaster recovery solution for your organization.
- Maintain constant email access for users round the clock, even if your Exchange server is not available.
- Minimize the risk of data loss due to on-premise server failures.
- Protect your email from spam, viruses, phishing, and other email-borne threats with built-in message security features

As a cloud computing service, Google Message Continuity requires no costly hardware or appliances, helping produce a low total cost of ownership. The service is easily managed through a simple web interface and includes customizable notification and authentication settings to inform IT administrators during Exchange server failures. Since all inbound, outbound, and intradomain messages are constantly synchronized between Exchange and Gmail, Google Message Continuity also helps ensure that users can still access their email messages, contacts and calendars during server outages so that your business can keep operating.

What's more, users can enjoy Gmail's full functionality during Exchange outages, including 25 GB of email storage, mobile email and calendar sync, and chat. After the outage is resolved, received messages and changes made to Gmail are synced back to Exchange.

Google Message Continuity also includes a complete set of email security features, so you can provide industry-leading security for your email systems without installing expensive software or appliances. This helps you block spam, viruses, and other external threats before they reach your organization, helping protect proprietary information that must remain confidential.

Google Message Continuity can bring Google's record of reliability and high availability to your on-premise email environment – all at low cost with and with minimal deployment complexity.

For more information, please visit: www.google.com/postini