WHITE PAPER

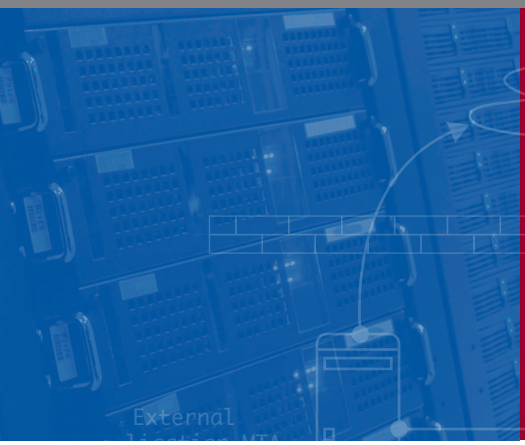

The Impact of Machine-Generated Messages on Enterprise Email Infrastructure

Gregory Shapiro, VP and CTO

SENDMAIL®

What is "machine-generated mail"?

In its most simple definition, "machine-generated mail" refers to mail submission by a non-human entity, such as a system process, embedded device, or cloud service. Another way to categorize machine-generated mail is to include all mail that is machine-to-person or machine-to-machine and exclude person-to-person and person-to-machine communication. Machine-generated mail comes in many forms, whether generated by customer care and notification applications, system monitoring utilities, scheduled batch jobs, meeting reminders, autoreplies, or a variety of other forms in which a message is generated in an automated fashion.

> "Sendmail estimates that at least 50% of the mail transiting a corporate mail infrastructure is machine-generated as opposed to human generated."

Based on enterprise customer data, Sendmail estimates that at least 50% of the mail transiting a corporate mail infrastructure is machine-generated as opposed to human generated, i.e., a person using a mail user agent (e.g., Thunderbird, Outlook, smart phone mail app, etc.). For the sake of this paper, we will also include mail generated as a result of human interaction, e.g., a human filling out a web form which triggers a mail message to be generated in the back-end or a person performing a banking ATM transaction and requesting an email receipt, as machine-generated mail.

What are the unique challenges created by machine-generated mail?

Message Volumes

Whereas a human may be able to produce a few messages per minute, applications can generate hundreds to thousands of messages per second. This brings about several challenges that are unique to machine-generated mail. At high volumes, safeguards need to be put in place to prevent "mail storming" of groupware or DMZ servers as well as external recipients.

Message Rates

Due to the lack of human intervention, without controls in place, anomalies can go undetected. For example, an application that was previously sending 1000 messages/hour suddenly sends 15000 messages/hour. If this was not expected, it could result in a loss in reputation if sent to external recipients (e.g., customers or prospects) or a self-inflicted denial of service (i.e., "mail storming") if sent to internal recipients. Monitoring is also needed to detect message transmission failure to prevent either mail loss for apps that don't have queuing abilities or deep queues that will impact application performance.

Policy Processing

Beyond rate limiting and message rate anomaly detection, machine-generated mail will have additional demands on the policy processing portion of the mail infrastructure. Machine-generated mail may have different policy and archiving requirements based on the application generating the mail and the content/purpose of that mail. This calls for a policy engine that can act based on identification of the machine-generated mail to perform policy based routing and regulatory and governance control on high volume machine-generated mail flows. Also, applications generating a high volume of automated template-based mail may be able to leverage techniques to enhance performance in hygiene filtering such as pre-expansion hygiene checks on the template instead of the generated individual messages.

> "It may be cost prohibitive or technically infeasible to migrate applications and devices that produce machine-generated mail into a remote cloud-based infrastructure."

Cloud Email

It may be cost prohibitive or technically infeasible to migrate applications and devices that produce machine-generated mail into a remote cloud-based infrastructure. Whereas it is relatively straightforward to automate the process of changing end-user mail configuration to point to new cloud-based mail infrastructure, in the case of machine-generated mail, each application needs to be

tracked to locate both the owner of that application as well as the server(s) running the application. Once that list is generated, each application on each server will need to be reconfigured or rewritten (in the case of hard coded mail configuration) to use cloud-based mail services. This can be challenging, as many of these applications may not be written to adhere to the mail protocol standards that cloud-based mail infrastructure services will require. Examples include the ability to handle MX lookups, handle multiple IP address records, output full and standards compliant headers, perform proper masquerading and cleaning of internal network information in headers, etc. Similarly, many cloud service providers place limits on mail flow or contents that will hinder the application's ability to deliver. Examples include message size, recipient count, message frequency, and daily aggregate count limitations. Adapting the applications can require re-engineering their messaging code (assuming you have the source code available).

Likewise, some machine-generated email applications/ devices may not have the ability to connect to a remote email infrastructure. They may not have network access or permission to connect to the Internet. They may not operate well with a high latency communication channel when talking to remote servers or may not be able to cope with unavailability, timeouts, or temporary failures/ protocol replies from a remote infrastructure. In some cases, it simply might not make sense to send machine-to-machine mail out to an external infrastructure, simply to be returned into the same site for use by another application.

"Organizations often have complex internal routing requirements necessitating an internal email backbone supplied by vendors such as Sendmail."

— Gartner

Email Is a Commodity and Other Fairy Tales

Gartner's Email Is a Commodity and Other Fairy Tales, February 2011 report states, *"...message-based workflow, alerting and notification services are all routinely integrated into email infrastructures, creating a complex overlay over the base messaging system. Simple SMTP integration can be transferred to the cloud, but tighter integration via proprietary APIs to the mail system are generally not transferable. Organizations often have complex internal routing requirements necessitating an internal email backbone supplied by vendors such as Sendmail."*

The final challenge related to interfacing machine-generated mail applications to the cloud is determining if you are comfortable placing the corporate data found in the machine-generated mail into the cloud without first having policy applied on-premises. To illustrate, there is some portion of your corporate data that is sensitive enough that the data itself cannot be stored externally for regulatory or corporate governance reasons (e.g., patient or financial systems of record). Machine-generated mail applications that make use of this data have the potential to leak that data. Therefore, any outbound mail from these applications must have data leak prevention policy applied. For example, a policy to make sure a message doesn't contain a list of names and social security numbers. This also implies that any machine-generated mail applications that make use of those highly sensitive data cannot themselves be migrated to the cloud without opening a data conduit between the remote cloud provider and the internal data source and therefore will require on-premises mail infrastructure.

"In some cases, it simply might not make sense to send machine-to-machine mail out to an external infrastructure, simply to be returned into the same site for use by another application."

What do I need to do to meet these challenges?

All but the smallest of businesses have many sources of machine-generated mail. Larger corporations can have hundreds of applications running on thousands of servers generating mail on a daily basis. In order to understand how the challenges presented may impact your organization, the first step is to identify the applications and other non-human sources generating mail and their business owner. Using that list, locate the servers or devices on which they run and build up an understanding of the requirements and message processing policies needed for each (including data requirements, archiving needs, data retention policy, and message rates). The goal of this effort is to identify, register, control, and monitor these machine-generated mail applications to reduce security risks, and protect brand reputation.

For those organizations already in the cloud or considering a cloud migration, determine if each can be migrated to the cloud. For those applications that cannot themselves be migrated into the cloud (e.g., because they need access to internal, on-premises data stores; require encryption at the gateway; need special mail routing; etc.), determine if the application/device can be reconfigured to communicate with a cloud-based mail infrastructure and if the cloud infrastructure can handle the requirements of that application/device. Finally, with the remaining applications and devices, determine the properly sized and minimum set of requirements for your on-premises mail infrastructure. Machine-generated mail applications that remain on-premises will, in most cases, require the on-premises infrastructure to handle rate limiting, policy control, intra- and extra-infrastructure routing, email standards adherence, and mail queuing. That on-premises mail infrastructure can then be integrated to work efficiently with the cloud infrastructure to create a balanced, efficient hybrid mail infrastructure.

To help kick start the first step of identification, the Appendix below contains a list of common services provided by machine-generated mail applications and devices. Use it as a starting point for finding applications and devices within your organization.

Conclusion

Every organization, large or small, has a number of applications and devices producing machine-generated mail. Unlike person-to-person communication, machine-generated mail has a unique set of challenges and requirements that must be understood in order to build a successful mail infrastructure, whether on-premises or remotely in the cloud. As you move through the process and start the planning for cloud migration or integration, to mitigate the risks, it is important to first modernize your existing mail infrastructure: chose which parts of that infrastructure to move and build out a plan to integrate the remaining on-premises infrastructure with the cloud-based infrastructure to create the aforementioned hybrid messaging infrastructure.

Additional Reference Material

- Important Information for Enterprises Moving Email to the Cloud Featuring Research from Gartner
  http://imagesrv.gartner.com/media-products/pdf/sendmail/issue1/issue1.pdf
- Can an Enterprise Email Backbone Infrastructure be Moved to the Cloud?
  http://sendmail.com/pdfs/whitepapers/Sendmail_WP_Backbone_to_Cloud_Field.pdf
- Moving to the Cloud: Important Things to Consider Before Migrating Your Messaging Infrastructure to the Cloud
  http://sendmail.com/pdfs/whitepapers/wp_moving_to_the_cloud.pdf
- *www.sendmail.com/resources/collateral*

Appendix: Examples of machine-generated mail

Machine or Application Type	Description
ERP Systems	
Record change notification	Triggers can be created to notify one or more people when an ERP record is changed in a material way
Reporting	SAP and other ERP system support automated, scheduled reports to be sent to one or more people using SMTP mail messages
New employee creation alerts, automated messages into other systems for account creation	HR systems can automatically send mail to managers and other work order systems to handle account creation and provisioning for new employees as well as account termination for departing employees
Notifications for credentials, password changes	When an end user forgets their password or makes a significant change to their account, mail is sent to the user to provide for password recovery or to notify the user of the change in case it was done without the user's knowledge
Ownership reminders and audits	Resources are owned or managed by various people within the company. Examples of resources include mailing lists, applications, servers, etc. It's common to send periodic reminders or audits to owners to verify the owner still exists and is still aware of their responsibilities
Employee communications	While every company has static mailing lists (potentially with directory driven membership), ERP systems can also create on-the-fly distribution lists and personalized messages to groups of employees based on different criteria For example, this allows an HR manager to send mail regarding a health plan change to only those employees subscribed to that health plan
Expense report processing	When expense reports are submitted online, the systems can support either notification of status (e.g., a message to a supervisor that an expense report is pending or to an employee that their expenses have been approved or rejected) or allow for an email based submission and approval system
Machine or Application Type	**Description**
Time off/time tracking	Online time tracking/vacation systems use email to notify supervisors of pending approval requests and time card exceptions as well as notify employees of approval or denial of requests
Backend for web-enabled applications	Any web-enable applications (e.g., BEA Web Logic web portals) can use SMTP email as their transport system by converting web form submissions into formatted email requests either to a human or to another application. Examples include customer service requests, RMA requests, etc.
Automated messaging from ERP applications	As of SAP Web Application Server 6.20, faxes and text messages (pager/SMS) can also be exchanged via SMTP. Any of the machine-generated mail discussed above can be packaged into an email message and sent to a fax or SMS gateway using SMTP as the transport. When something goes wrong with the automated processes (e.g., order taking), notifications are sent
SMTP email used a protocol between multiple automated systems	Procurement system Ariba sends mail to special alias for CRM for "email to case" (e.g., setup support portal for new customer) and "email to sales order" functionality Reports on queued orders between systems are emailed as a form of monitoring and alerting
CRM Applications	
Send quotes, license keys, support reminders to customers	CRM systems support automatic mailings to customers for quotes, license keys, expiring support/license reminders, etc.
Bidirectional support ticket updates via email	While the support case actually lives on the CRM system, customers submit cases via email, which automatically creates a new case in the system and responds with a case ID. From then on, both customers and support technicians can use email through the CRM system to keep in contact and create the case records/audit log
Lead notification to sales team	When new leads are found, either through data loading from an event, web form submission, or added by sales development representatives, a notification is sent to the territory's sales team
Order tracking	Although more common in the SMB market, customer notifications of order status is sent via email. Likewise, sales order and shipment reports can be scheduled for email delivery to the appropriate people within the company
Telephony Voice Mail Systems	
Voice mail converted to audio files & sent as email attachment	Modern phone switches can not only save voice mail to the system for retrieval via phone but can also convert those to an audio attachment and email the voice mail
Call reports	Voice mail systems can mail scheduled reports to administrators as well as missed calls to recipients (similar to the voice mail email)
Bidirectional SMS/MMS gateway via email	Email to SMS/MMS (and vice versa) gateways allow companies to communicate across messaging systems. These are usually used for alerting purposes

System Alerts/IT	
SNMP/Nagios alerts via email and pager email gateway	Monitoring software such as Nagios support email notification when a monitoring threshold is triggered. This email can be to a human or to an SMS gateway
Application notification of events (e.g., bug tracking systems, source code control changes)	Many applications include email notification of certain type of events
Tripwire and other security scans	Security scanners such as Tripwire run either in a monitoring mode or in a scheduled mode to scan for anomalies and alert administrators via email
Nightly system status notifications	Most UNIX based systems run rightly checks on usage, disk status, important log messages, etc. The results of those checks are mailed to the system's administrator(s)
Build reports	Results for automated build and test systems are mailed out to the developers responsible for the code being built/tested
Cron reports	UNIX systems include the ability to schedule automated processes via cron. Any output from those automated processes is mailed to user who created the cron job
Log watcher reports	Log watching/reporting products, such as Splunk, allow scheduled reports to be generated and mailed. They also support the ability to apply special searches to logs and set conditions which, when triggered, generate an email alert
Backup reports	Backup systems send status reports of the backup, what was backed up, which tapes need to be changed, etc.
Reminders (calendar, certificate expiration, domain renewal)	IT staff typically maintain a calendar of upcoming events such as certificate expiration, domain expiration, etc. Those calendar events trigger an email reminder to the IT staff to take action
Non-delivery reports (e.g., bounces, postmaster notifications)	The mail system itself generates automated messages to IT staff when there are delivery problems that need to be addressed by those maintaining the mail system. The most common of these are MTA postmaster notifications (aka, non-delivery reports)
Machine or Application Type	**Description**
Office Machines	
Copier scan to PDF & email	Modern copiers can scan documents to PDF files, which are either deposited on a shared file server or sent to one or more individuals via email
Fax machine convert to PDF & email	Like the modern copiers, modern fax machines can convert faxes being received to PDF and email them instead of printing them to paper. In some cases, this is the same physical machine as the copier (i.e., multifunction printers)
Badging system alerts & report	The badging system for entry access can send alerts and reports to facilities managers. The badging systems may interact with other systems to provision new employees and remove departed employees as described in the ERP section

Sendmail Messaging Architecture Review

To help enterprises properly plan an email infrastructure cloud migration project, Sendmail offers a comprehensive Messaging Architecture Review. A Messaging Architecture Review is comprised of a thorough review and assessment involving input from your company's messaging team and other key business units. To support this, Sendmail Messaging Architects review your existing architecture and implementation, as well as current issues and concerns, to develop the following:

- Business objectives
- Short-term recommendations
- Long-term recommendations
- Recommended roadmap
- A comprehensive report and presentation, delivered to the organization for review and analysis.

About Sendmail, Inc

Sendmail provides enterprises with global email connectivity, routing, and message delivery between people, systems, and applications. The Sentrion Email Infrastructure Platform reduces enterprise IT complexity and lowers the cost of integrating critical message content with business processes. Since 1982, thousands of enterprises around the world have relied on Sendmail open source and the award winning Sentrion hard and virtual appliances, and its suite of applications, for intelligent email backbone and secure message gateway infrastructure. Unlike proprietary email-security appliances, Sentrion provides full-content inspection and policy-based delivery for all collaborative groupware and machine-generated email, whether deployed on-premises, in-cloud, or via mobile devices. Sendmail is headquartered in Emeryville, CA with sales and support offices throughout the Americas, Europe, and Asia.

Sendmail, Inc.
6475 Christie Avenue, Suite 350, Emeryville, CA 94608 USA

Tel: +1-888-594-3150 | Fax: +1-510-594-5429
Email: info@sendmail.com
www.sendmail.com