

The Six Steps to Reducing Business Latencies

An In-Depth Examination of the Inefficiencies
of 6 Common Business Processes

Vineyardsoft Corporation
www.vineyardsoft.com
800-850-8055

Introduction

Efficiency. When you look at all the ways that an organization tries to improve their business, they all come down to a single goal. Efficiency. The investment in the most recent technologies such as phone systems, computers, and computer software is done with the explicit hope these investments will increase organizational efficiency. Similarly, investing in training is done with the expectation that the training will result in a smarter staff, one that makes fewer mistakes and takes less time to do routine tasks.

Although these two investments – technology and training – are quite different, they both have the same ultimate goal: **Efficiency**. In other words, *“the ability to do one’s job faster and smarter than before”*.

Faster and Smarter. Almost all corporate initiatives can trace their desired results to either (or both) of these goals. But “Faster and Smarter” is becoming more difficult for businesses to achieve.

In the “old days” (1970s thru the 90s), it was easy. Businesses looked **outside** of their walls, in search of **fast technologies** to support their internal processes. Manual processes were replaced by automation, and “old” automation was replaced by newer and faster automation.

This is no longer the case. Today’s companies already have fast technologies in place. Investing in a newer technology that processes data one millisecond faster than an old technology is not the answer.

Companies must stop looking **outside** of their own environment, in search of **fast technologies**. Instead, companies must begin to look **inside** their own business, in search of **slow processes** that can be improved.

These slow processes are referred to as a company’s **Latencies**.

Once Data Latencies are identified, organizations can begin to implement processes (and technologies) specifically designed to address them. This paper will detail the six most common areas for Latency within an organization, and will explain not only how they can be identified, but the processes and methods required to reduce them.

The 6 Business Latencies

Although every business has its own unique set of business processes and challenges, it is possible to identify six “generic” areas of business activity that typically contain Latencies. These areas are:

1. Executing Redundant Tasks
2. Identifying and Responding to “Trouble Spots” (Exception Management)
3. Profiling & Loyalty-Building
4. Monitoring and Processing Incoming Email
5. Generating and Distributing of Reports
6. Integrating Data Analysis & Response

The following pages will detail each of these in the following terms:

- ➔ Definition of the Latency Area
- ➔ Real-Life Latency Examples
- ➔ The Latencies’ Impact on Business
- ➔ How to Reduce the Latencies

Note: Zero-Latency Technologies

Although this paper will discuss examples of business scenarios with embedded Latencies, the methods by which these Latencies can be reduced, and the functional requirements of Automation to address these Latencies, this paper is not designed to discuss or review commercial Zero-Latency technologies on the market today. These technologies are often categorized under the heading of “Business Activity Monitoring” (or “BAM”) solutions and range in price from \$3,000 to \$300,000.

To explore the appropriate Zero-Latency technologies for your organization, it is suggested that you determine which of the six Latency-prone areas are evident within your organization and then evaluate each potential solution in regard to its ability to address those areas of need.

Latency #1: Executing Redundant Tasks

Definition: The process of performing those repetitive tasks which include some amount of human intervention.

It’s important to note that a redundant task in itself does not necessarily include Latency; only if the redundant task requires some amount of human intervention does Latency enter the picture to a greater or lesser degree.

Examples:

The following are examples of redundant task activities that require some degree of user intervention. Note that by definition these examples are not “exceptions” to business-as-usual; they each represent a recurring activity that requires one or more people to expend time and effort that could otherwise be used elsewhere.

- A salesperson who generates and delivers “standard” messages such as Order Confirmations, Quotes, Expirations Notices, and so on.
- A person in Finance who generates and delivers Invoices, Payment Reminders, Dunning Notices, etc.
- Support Reps who, after they close an open call, draft and send a “resolution e-mail” message back to the client who reported the problem.
- A person in Shipping who tells someone in Sales that the delivery of a specific item has been delayed.
- Marketing staff, among whose job it is to monitor visitors to your website and see if any visitors have made special requests or are from specific companies.
- IT staff, whose daily tasks include looking for specific “import” files, notifying the appropriate people about these files, and processing them as needed.

Business Impacts:

Any time you allocate an Employee to do work that could be accomplished by Automation, you are wasting precious business time. It’s a plain fact that automation works faster than humans, and so a five-minute human task can often be accomplished in five seconds (or less) by automation. Multiply that one task by the number of times that it’s executed on a daily basis, and you begin realize just how much time you’re wasting per day on just that one task.

If the task can be done as capably by Automation as by an Employee, there’s no reason why an employee should be assigned that task.

And then there is the concept of Human Time, most appropriately summed up in the business refrain:

“Don’t you have something better to do?”

Not all human tasks can (or should) be done by Automation. But this doesn’t mean that you shouldn’t explore every opportunity to see where Automation is a good fit. After all, “Human Time” – insofar as our unique abilities to converse, brainstorm, and problem solve, are all very valuable commodities, and we should take every opportunity we have to maximize the available time to use these skills, and (correspondingly) minimize those times when such skills are not required.

Anytime you see a repetitive activity taking place around your organization, your first thought should be “Is there some way we can automate that?” Too often organizations assume that just because a repetitive activity has always been done by a human, it must continue to be so. And chief among the reasons why this assumption is made is that organizations incorrectly assume that to remove the human from a redundant task means also to remove the “human element”.

In some cases this may be true, such as the necessary presence of human beings (and not computers) to staff most organizations’ Customer Support Departments. There, the human element (albeit often combined with technology) is usually essential.

But consider the example of the salesperson who manually generates (and delivers) each client’s Order Confirmation. What is the “human element” in such a task? The human element in this scenario is the friendly, personalized Order Confirmation message. Is such a personalized message beyond the means of today’s technology? Not at all; this is an ideal task for Automation, and it enables employees to devote more time to more demanding (and profitable) endeavors.

The Latencies connected with Redundant Task Execution also include **Data Integrity**. Although humans have certain traits that distinguish us from (and give us unique advantages over) Automation, we too have our weaknesses. Chief among them is our proclivity to make mistakes. Consistency of performance is simply not one area where humans outperform Automation.

And when you consider that the overriding goal of reducing Latencies is improved organizational efficiency, the one sure-fire way to negatively impact efficiency is to make mistakes. Mistakes themselves not only waste time, but more often than not end up creating ancillary business situations which themselves consume additional resources. If for no other reason than Data Integrity, the Automation of Redundant Tasks is key to an organization’s efforts to reduce Data Latency.

Reducing the Latency of Executing Redundant Tasks:

So – how does an organization go about reducing the Latencies caused by Redundant Task Execution? The simple answer is “automate” – but consideration and planning need to go into that decision.

First of all, determine if a redundant task is a good one to automate. Consider the “human element”, and decide whether the amount of time expended on the task justifies the investment in Automation.

Second, consider the type of functionality required by suitable Automation. Although some recurring tasks can be automated through Job Scheduling Automation, such Automation would **not** be sufficient to address the examples listed on the preceding page.

Why? Because these recurring tasks have one special requirement – their execution is conditional. Whether it’s the presence of a newly-entered order, or the submission of a web form from a specific client, the execution of these redundant tasks is based on the presence of certain conditions. Job Scheduling Automation (in general) does not have the ability to “detect and trigger” when tasks should be performed.

Detection in particular is key. Most organizations fail to automate redundant tasks simply because they believe Automation is not capable of identifying the unique set of conditions required to trigger a task.

This is not true; *intelligent* detection and response Automation solutions do exist, and they are indeed the answer to this particular problem.

Latency #2: Identifying & Responding to “Trouble Spots”

Definition: The ability to automatically perform periodic checks of business data to identify and respond to existing conditions of data which indicate a problem or potential problem. This is often referred to as Exception Management.

This is one of the more challenging Latencies for a Business to recognize, since most organizations’ approach to trouble (or potential trouble) is to “see it, fix it, and forget about it”. As such, the identification of recurring trouble spots (which is where Latency is most apparent) requires an in-depth analysis of an organization’s business processes and asking such questions as:

“What happens in our business when . . . ”.

Although it’s human nature to fix a problem and move on as quickly as possible, the exercise of working backwards from the problem (to explore why it occurred and how to prevent it) allows us to identify the true cause of the Latency.

Examples:

The identification of (and response to) trouble spots in a business is commonly referred to as Exception Management. This term represents an approach to business whereby the norm is accepted, and only evidence of processes that stray from the norm are identified and acted upon. In that sense, Latencies are easy to identify, as you start out by identifying a specific process or rule and simply ask yourself whether it’s feasible for that process or rule **not** to be followed. For example:

- ➔ Rule: Prospective clients must be contacted at least once a week. Latency: A prospective client has not been contacted in three weeks.
- ➔ Rule: Inventory levels on certain items must not be let to drop within 10% of their re-order level without a PO being issued. Latency: An item drops to within 5% of its re-order level and does not trigger any kind of action.
- ➔ Rule: The completion percent of a project should never vary more than 20% from the percentage of budget used up. Latency: A project is 60% done and has used 95% of its budget.
- ➔ Rule: An individual Support Rep should not have more than 6 “critical” tickets assigned to them. Latency: A rep has 14 “critical” tickets assigned to them.
- ➔ Rule: Sales reps are required to synchronize their local database with the corporate DB every 4 hours. Latency: Nobody in Sales has synch’ed within the last 3 days.

-
- Rule: If a delivery is delayed beyond the Required Ship Date, the client must be notified. Latency: An order that was due to ship last week still has not shipped and the client has not been informed.

Business Impacts:

Back in the 1980s, a company that sold System Monitoring software came out with an advertisement that showed an IT Director wearing a Fireman's helmet and sitting in the seat of a child's Fire truck. The ad was a huge success, and the response from every IT manager was the same: *"I know what it's like to put out fires."*

Reducing Business Latency isn't about putting out fires; it's about preventing them from happening in the first place. Once the fire has occurred, so too has the Latency; time has been lost starting from the very instant that the business rule or process was broken. And the Latency doesn't stop once the problem is located; correcting the situation always requires additional time and effort – time and effort that could be better spent elsewhere. It's a simple rule of thumb; the more time that goes by between the initial breakage of a business rule or process and the detection of that breakage, the greater the Latency.

But negatively impacting an organization's meaningful working hours usually isn't the worst fallout from an inability to pro-actively identify trouble spots. Worse still is the damage that this Latency has on Customer Relations.

Studies have shown that it takes an average of eight to twelve positive interactions with a customer to earn their loyalty, but only one negative interaction to lose it. With those kinds of odds, it's imperative that you do whatever you can to minimize the number of negative interactions. Once a negative interaction has occurred, the best that one can hope for is that the customer will not stop doing business with you, and will not spread the word of their negative experience. At worst, a negative interaction can cost your company future sales from this client, and from any other clients who hear about the negative experience.

Surprisingly enough, there is the potential for a silver lining to a negative customer interaction.

In a recent survey of customers who have continued to be loyal to certain organizations – in spite of having experienced one or more negative interactions with them – the survey results showed that clients place an extremely high value on how an organization responds to a problem when it occurs. As one customer put it: "The surest way for a company to gain my loyalty is to show me how well they respond to a problem that I'm having. It's easy for a company look good when everything is going as planned; but it's how a company responds to a problem that shows what they're really made of."

And problems will occur. In spite of the best technologies, processes, and training, problems happen and it's within an organization's power to turn that negative interaction into an Opportunity to show just how responsive they are to a client's needs. But the Opportunity is a delicate one. If too much time goes by before the problem is detected and responded to, no amount of good service is going to turn that into a positive experience. Likewise, if the responder to a problem does not have the appropriate information to help the client, the Opportunity is similarly lost.

Speed of detection and **access to meaningful data** are key components to reducing the Latency from these types of encounters.

Reducing the Latency of Exception Management:

Performing Exception Management (identifying and responding to potential trouble spots in your business) can be an extremely time-consuming process if done manually. But to automate Exception Management is no simple task. And the most difficult part of that task is to identify those business conditions that fall under the heading of "Exceptions".

One effective way to approach the identification of business exceptions is to choose a process (such as a sales process, a customer support process, or a manufacturing process) and analyze the process to see if you can apply some or all of the following four tests:

- ➔ **Date / Time Sensitivity.** Does the process have dates and/or times associated with it? If so, are there scenarios whereby an item going through this process could have dates or times that are indicative of a problem? (E.g., an open sale whose forecast close date is in the past.)
- ➔ **Numeric Thresholds.** Does the process have numeric values associated with it? If so, could an item have one or more values indicating that something is amiss? (E.g., a high priority support call that has been open for more than 4 hours.)
- ➔ **Too Many / Too Few.** Is it possible that too many (or too few) of an item are in a problematic state? (E.g., too many overdue activities assigned to one employee.)
- ➔ **"Special" Objects.** Are there specific items (or types of items) that should be handled differently from the rest? (E.g., a support call from your largest client.)

Once you are able to identify these Exceptions, you can begin to explore Automation that reduces the Latency of identifying and responding to these exceptional conditions. In general, three components are typically required in a solution that automates Exception Management:

1. **Exception Monitoring.** The ability to periodically and automatically monitor business data to see if "exception" conditions exist
2. **Exception Notifications.** The ability to notify select individuals about these conditions
3. **Exception Workflow.** The ability to automate the execution of one or more business processes ("response actions") that occur when an exception condition is detected

The last of these (the ability to *respond* to an exception condition) is worthy of extra consideration. This function raises the capability (and commensurate benefits) of an Automation solution insofar as it enables an organization to create a system of automated responses to their exception conditions. And although it is true that the act of notifying an employee is also a type of "response action", it is still incumbent on that employee to take some action of their own. An Automation solution that both alerts an employee **and** takes some action (on behalf of that employee) provides a significant advantage.

For some exception conditions a simple Alert is sufficient; but for other conditions, the need to perform an action – whether scheduling a follow-up activity, creating a Support ticket, et cetera, is essential to their goal of reducing their business Latencies.

The bottom line is that if you do decide to automate the Exception Management process, make sure that you know:

1. What types of conditions you need to monitor (record-level conditions, aggregate or calculated conditions, the ability to detect when key business data has changed, etc.)
2. How you wish to notify people (email, pager, cellular phone, fax, etc.)
3. What kind of automated responses you require (the ability to trigger updates into your business applications)

Based on the answers to these three questions, you will be able to easily identify which Automation Technologies are right for your organization.

Latency #3: Profiling & Loyalty-Building

Definition: The ability to dynamically analyze current and historical business data to identify potential sales opportunities, and automatically initiate tailored marketing programs and loyalty-building messages.

This Latency is often overlooked entirely, as far too many organizations fail to look towards their existing client base as a source for incremental revenue. This Latency needs to be viewed from two perspectives. First, this Latency must be considered from the perspective of an organization's ability to identify client profiles that indicate incremental sales opportunities. And second, this Latency must be considered from the perspective of an organization's ability to build Customer Loyalty through a series of meaningful, pro-active customer interactions.

Examples:

Most organizations do not expend much time or effort on either client profiling or loyalty-building. And the few organizations that do try to do so end up spending significant amounts of time and energy for little in return. The following business processes, although well-intentioned, are so imbued with Latencies as to make them impractical for any organization to implement.

- A Sales Manager meets with their staff at the end of every month to review what clients have purchased what products, and who would be a good target for a new product or service offering.

-
- A Support Manager reviews weekly reports to see which customers have called into Support about what issues, and how often. The goal is to identify those clients in need of Training, and those clients who might be interested in a particular Service Plan.
 - A CFO reviews clients' Receivables balances to see which clients average open balance is within a certain percentage of their credit limit, so as to offer that client a variety of different payment options.
 - Salesreps manually generate and send courtesy "thank you for ordering" letters to each customer exactly 30 days after they place an order.
 - The Accounting department manually drafts e-mail messages to clients reminding them that if they pay within 'x' days they are eligible for a discount of 'y' percent.

Business Impacts:

Why are the preceding business processes, unquestionably well-intentioned, so filled with Latencies?

Two reasons:

1. They are tasks intensive in manual labor
2. They do not happen in "real-time"

This paper has already discussed the value of "Human Time" to an organization; thus it is not difficult to see what a poor use of time it is to assign one or more employees to wade through mountains of data in the hope of finding clients who match certain profiles.

And what if those profiles are eventually found? What will be the single most important factor to determine whether an organization will be able to beat their competition to this prospective sale?

Timing. The organization that waits until the end of the day, end of the week, or end of the month to analyze their data and identify potential prospects will always lose to the organization that is able to do the same analysis – but on an on-going, dynamic basis.

It is a fact that the vast majority of those organizations who choose to analyze their client base do so manually and infrequently. Analysis often occurs days, weeks, or even months later than it should. And in spite of the fact that most organizations readily agree that potential problems need to be identified and acted on as soon as possible, those very same organizations do not recognize the same need with potential opportunities.

And, similar to potential problems whose Latency and impact grow with the longer they remain unidentified and unresponded to, the likelihood of closing a potential opportunity diminishes as the amount of time it remains unnoticed increases.

But real-time access to potential opportunity data is only half of the answer. The expense in manual labor ("Human Time") is an equally formidable roadblock in the attempt to profile existing clients. The prospect of generating piles of reports for the sole purpose of identifying potential opportunities is high on effort and low on appeal. And when you combine the manual effort required, with the additional Latency that comes from putting off

such a project until it's absolutely inevitable, the chance of finding and closing a potential sale is slim indeed.

The obvious impact on an organization's business by choosing either not to profile clients, or not to profile them in an efficient manner, is lost sales. But there's more to it than that.

There's the recurring subject of Customer Loyalty. Loyalty is built on many things, not the least of which is the ability to instill in a customer the feeling that you – as their Vendor – care about them, and are looking after their best interests. And there's no better way to demonstrate that than to follow-up with them in an intelligent (“interested”) fashion and to offer them products, services, tips, etc. that relate to what they are doing. This confirms that you know them, care about them, and have their best interests in mind.

This type of interaction is often overlooked, as it is very difficult to measure the impact of this type of interaction on Customer Loyalty. It's also difficult to measure the corresponding impact of Customer Loyalty on organizational revenues. But periodic customer “touches”, whether they take the form of an order “thank you” letter, a reminder about special offers or discounts, or simply a “Happy Birthday” greeting, all contribute towards building a closer, and more loyal relationship between your organization and your clients. Without a system that supports such Loyalty-building actions, your organization joins the ranks of “all of the other” companies that your clients do business with.

Reducing the Latency in Profiling & Loyalty-Building:

Instituting a Profiling and Loyalty-Building system is not high priority for most organizations. As a result, it usually manages never to get done.

There are three main reasons for this. First, an organization has to consider what **type** of profiles they wish to identify. This is a time-consuming and difficult process, as the task of simply arriving at a consensus as to what those profiles should look is often a challenging one. Secondly, organizations typically are reluctant to spend the **time** and **resources** to analyzing customer data in the hope that they will uncover one or more clients who meet the profiling criteria. And thirdly, by virtue of the time and effort required to search for likely profiles, the whole profiling task is usually accomplished so many days or weeks **after** the profile first emerged that the resulting window of opportunity to act on the profile has long since closed.

The key to jump-starting this process is to divide it into small, easily-accomplishable tasks. For example, an organization should focus (initially) on the easily accomplishable task of winning **Loyalty** from clients instead of immediately trying to win more sales from them.

(In other words, do not, for example, start a Profiling initiative by trying to identify those clients who are in industry '**a**', have purchased more than '**b**' dollars of products '**c**' and '**d**' between dates '**x**' and '**y**' and therefore might be a good candidate for product '**z**'.)

You can focus on winning Customer Loyalty with some very simple – and very “considerate” events. Examples include “thank you for your order” messages, reminders about available discounts, birthday greetings, 30-day order follow-up messages, and so on. Not only are

these types of tasks very easy to automate, they have the distinct advantage of gaining you Customer Loyalty with very little time, money, or material expenditure on the part of your organization.

Tasks that are designed to gain Customer Loyalty will also often gain you unexpected sales opportunities (using a “soft sell” approach) than a more direct “you purchased this and so might be interested in that” approach. Loyalty has that affect on people. And over time, you can begin to complement your Loyalty-Building tasks with those that are a little more sales-related. The combination of the two are a sure winner.

As for the Automation required to support both Loyalty Building and Profiling related tasks, the Automation needs to balance Data Mining capabilities (to determine if clients meet certain criteria) with Flexible Notification capabilities (to personalize outgoing messages).

And there are two additional capabilities that you should keep in mind when considering the automation of Loyalty or Profiling Building tasks. The first is the ability of an Automation solution to keep your Customer Database(s) informed about the actions that the Automation solution has taken. For example, if an Automation solution sends a birthday greeting to a specific client, it’s important that the Automation solution note that delivery within the corresponding Customer Database(s). Similarly, if, for example, you want a client’s birthday to result in a “Happy Birthday phone call” from their salesperson, an Automation solution should have the ability to schedule that call in an underlying business application.

Second – and very important – is the consideration of what happens when and if the customer, upon receiving an Automated message, responds to that with a message of their own. One of the greatest contributing factors to the **loss** of Customer Loyalty is the inability for a company to respond quickly and intelligently to a customer’s request or inquiry. So – whenever you consider implementing a process to improve interactions with your clients, keep in mind that those communications are a **two-way street**. Without the ability to **quickly and intelligently respond to your clients**, chances are good that your clients won’t be in the mood to pay attention to your messages in the first place.

Latency #4: Monitoring & Processing Incoming Email

Definition: The ability to analyze the content of incoming email messages, determine if a message meets certain criteria, auto-route the message to the appropriate recipients, auto-reply to the original sender, update Customer Information databases with details from the incoming message, and schedule intelligent follow-up activities within those same databases.

Note that this definition applies both to email that is sent directly to an organization, as well as email that is created from intermediary sources, such as from an organization’s website. (Many websites utilize “web forms” which allow a visitor to submit information to an online “form”. This information is then loaded into an email message which is automatically sent on to a specific email account within an organization.)

Examples:

- Emails sent to support@yourcompany.com. Allowing customers to log their own support questions with an organization via email is good. Having no way to

automatically review, categorize, and log these messages into a support database is bad.

- Emails sent to info@yourcompany.com. Allowing customers to send emails to “generic” departments within an organization is good. Having no way to automatically interrogate and re-direct those messages to the appropriate person(s) is bad.
- Emails sent from specific Clients. Having the ability to identify email from specific “key” clients is good. Having no way to automatically create and assign follow-up tasks from those messages is bad.
- Web Forms Requests. Allowing clients to use your website to make requests (such as for information or assistance) is good. Having no way to auto-respond to those requests (such as automatically sending back requested information) is bad.
- Web-based Enrollments. Allowing clients to use on-line enrollment forms to sign up for classes is good. Relying on manual intervention to update underlying databases with latest enrollment details is bad.

Business Impacts:

Nowhere today is the phrase “**Lost Business**” so appropriate as when applied to an organization’s management of their incoming business-related email. Business-related email is often misplaced, ignored, or simply deleted. And the business that would have been gained from that email is lost right along with it.

Conservative estimates put the loss of meaningful e-mail at almost 25%. That’s not to say that one out of every four messages is actually misplaced, never to be seen again, but rather that one out of every four messages is handled so poorly (either lost, never reviewed, mis-directed, or simply not acted upon in a timely fashion) that by the time it does come to light, it has lost all meaningful value. The lead is gone elsewhere, the sale does not happen, the client is irate. No matter how you categorize it, it’s a **loss**.

The real shame of “lost” emails is that unlike a potential sale whose loss is logged and tracked in some kind of application, there are few (if any) ways that you can track the negative impact caused by lost email messages. And such losses can have a profound impact, as emails affect every part of a company from sales (“*Where did that inquiry go?*”) to Customer Service (“*Has anyone seen the error message that John Smith emailed to us?*”), to Finance (“*Did Karen ever email us her credit card details?*”).

From a Latency perspective, the amount of “Human Time” that organizations devote to monitoring, tracking down, and processing incoming emails is staggering. If ever there was a Redundant Task (see **Latency #1**) that demands Automation, processing incoming email is it.

It’s bad enough that most organizations use manual processes to monitor, review, and respond to incoming email. For most employees, email monitoring is one of their many responsibilities. This means that the Latency time between the receipt of an original message and the execution of the necessary follow-up actions (such as sending alerts, updating databases, and sending an acknowledgement) is significant. We’re not talking

about Latency time in terms of **minutes** here; we're talking **hours** (if lucky), and more often than not, **days or even weeks**.

Even if it is the sole responsibility of one or more employees to manage incoming e-mail, there is the inevitable probability of human error to further delay and damage the processing. Simply put, using employees to process vast quantities of incoming email is very inefficient. There will always be some small percentage of email requiring human intervention to process correctly, but that is exactly what Automation is for – to process the large amount of “normal” data and to request human intervention only when needed.

The effective management of incoming email can also have a profound impact on the responsiveness of an organization to client needs, and, correspondingly, on an organization's competitiveness and capacity to retain (and build) Customer Loyalty.

Consider the client who sends email inquiries to three different companies. Company #1 responds to the email after 5 days. Company #2 responds after 5 hours and company #3 responds in 5 minutes. Not only will the client remember Company #3 for their responsiveness, Company #3 (by being first) gains the advantage of setting the client's expectations to most closely fit their offerings. Before even one bit of product information is distributed, Company #3 has a distinct advantage.

Rapid response is also a powerful ally in building Customer Loyalty. Customers will tolerate quite a bit from the companies they do business with – ***except for being ignored***. The organization that shows immediate response to a client's inquiry – even if that response is no more than an acknowledgement of receipt of that inquiry – is more likely to retain that client's business.

Reducing the Latency in Monitoring & Processing Incoming Email

Of all the Latencies discussed in this document, monitoring and processing incoming email is the most challenging, as there are four distinct steps in the email management process. Latencies can (and do) occur in all four steps.

But the news is not all bad. Unlike other Latencies where the choice is either to Automate or not to Automate, the very complexity of email management enables an organization to address its inherent Latencies in individual steps. Using Automation to address even some of the steps dramatically reduces the overall Latency and provides a significant improvement to an organization's overall business.

The four steps to effectively monitor and process incoming email are:

1. Analyzing the incoming message to see if it is of importance
2. Re-routing the message to the appropriate recipients and/or Alerting affected persons
3. Sending an acknowledgement (or “receipt confirmation”) back to the original sender
4. Updating the appropriate Client Databases with the details of the incoming message and (if required) scheduling appropriate follow-up activities.

Steps #1 and #2 are relatively simple to automate and should be done by every organization that utilizes “generic” email addresses, such as sales@yourcompany.com or support@yourcompany.com. Most commercial email systems (including Microsoft Outlook™) have the ability to auto-route messages based on criteria within the message itself. It could be as simple as routing all messages based upon the address they were sent to, or more complex routing that re-directs messages based on content within the message itself.

Even more sophisticated routing is possible. One example would be to check a sender’s e-mail address against a database of contact records to determine who that client’s “account manager” is. The email message would then be re-directed to that person. Routing such as this usually requires an additional piece of Automation to perform this comparison.

When considering the auto-routing of a message, consider also whether the receipt of a message should trigger other alerts as well. For example, the receipt of a support-related email message might necessitate the paging of a client’s salesperson. The salesrep doesn’t need (or want) all of the technical details of the message, but they certainly would appreciate a quick ring on their pager (or possibly cellular phone) to let them know that one of their clients is in need of assistance.

Automation is also the answer to steps #3 and #4. Step #3 (sending a confirmation message back to the sender) is one of those actions that can be set up quickly and easily, and yet has a profound impact on the perceived responsiveness of your organization.

Even if a confirmation or acknowledgement message says nothing more than “*Thank you for sending in your request*”, it tells the sender that:

- a) Their message has not be lost or sent into a black hole
- b) Their request is being worked on.

In truth, the second of these may not be entirely accurate, but the sender’s perception is that they have received a response back from you. That immediate response – if nothing else – buys you a little extra time before you need to begin working on and responding to the original request.

Step #4 (updating business applications with content from an email message and also updating one or more applications with intelligent follow-up activities) is the most challenging step to automate. If you are interested in using Automation to address this task, you will be best suited to look at “middleware” software, or (more specifically) a class of software called “Business Activity Monitoring” (often abbreviated to “BAM”).

Middleware is a type of Automation software that sits “in the middle” of your applications and fosters communication between them. By its nature, Middleware works with the widest variety of applications, and it usually includes embedded tools that let it combine, analyze, and exchange data equally well with all of them.

The “application independence” of middleware is key, because it gives an organization the flexibility to choose the best possible business applications for their needs, and then have a single Latency-addressing “middleware” application that works with all of them. This eliminates the need to learn the varied Latency technologies of multiple business applications, and gives an organization a robust single Latency-addressing middleware solution that will continue to work as an organization’s needs grow and change.

Latency #5: Generating & Distributing Reports

Definition: The ability to automatically trigger, generate, and deliver reports to the appropriate recipients.

Note: Report Distribution may be divided into two categories; **Scheduled Reports** (reports that are generated and distributed according to a recurring, periodic schedule), and **Triggered Reports** (reports that are generated and distributed only when certain business conditions exist).

Examples:

The following are all examples of reports that organizations typically require in their day-to-day business. If run manually (i.e., if it is an employee's responsibility to run and distribute the report), the process contains unnecessary Latencies.

- A daily "Open Support Calls" Report needs to be generated every morning at 9 AM and emailed to the Support Manager and Sales Manager.
- A daily "Overdue Activities" report needs to be generated and emailed every morning to each Salesperson who has overdue activities.
- Once a week, a Report should be run for any Salesperson whose pipeline contains less than \$50,000 in prospective sales. The report is emailed to the Sales Manager
- A daily "Project Management" report needs to be generated and emailed to over 100 employees within an organization.
- Once a week, any client who has more than \$10,000 in overdue Receivables should receive a summarized invoice that lists all overdue payments.
- Whenever a new order is entered into the system, an "Order Confirmation" report should be delivered to the corresponding customer.

Business Impacts:

The time when an organization had one or more people whose jobs included running countless reports for their Manager or Supervisor are (fortunately) mostly behind us. Automated Report Distribution applications (or add-on modules for commercial Reporting applications) are fairly widespread and, in general, are very affordable. Where standard Report Distribution applications fail, however, is in their inability to **self-determine** whether a report should be generated, and (if so), what data the report should look for, and whom the report should be sent to.

But let's begin by looking at the very first example from the previous list, the "Open Support Call" report. Clearly, the task of running such a report and delivering to the appropriate recipients is not a good use of "Human Time". Thus one would expect that most organizations needing to run a report like this would have its generation and distribution automated. Even so, there are still some hidden Latencies.

For example, what if there are ***no*** Open Support Calls? The result will typically be a ***blank report*** which is sent to the appropriate recipients. This is a minor annoyance and a waste of time – both on the part of the person who generated the blank report and for the recipients of that report.

But “Human Time” is not the only resource that is impacted by reporting Latencies; so too is System Processing time. Consider the “Project Management” report that needs to be emailed to 100 people. Is it really a good use of system resources to generate and e-mail 100 identical reports? How will those reports affect the processing speed of an organization’s email server? Reports are notoriously “resource-hungry”, and the effect of running (and/or emailing) 100 reports would certainly adversely affect system performance.

Report generation and distribution’s greatest Latency, however, is in its inability to be spontaneous. By definition, “Report Distribution” is governed by a ***schedule***; reports may be scheduled to run every couple of hours, once a day, or even every week. Unfortunately, the business conditions that warrant the creation of a report can occur ***anytime***.

For example, a Sales Manager might receive a report called “Overdue Sales Opportunities” every Monday at 9 AM. If an opportunity is due to close on a Tuesday – but does not – the Sales Manager will not know about it until next Monday; ***six days later***. By the time that overdue sale is identified, it’s probably too late to do anything about it.

Unless someone is actively watching business conditions, the conditions that ***warrant*** a report and the scheduled ***running*** of that report may occur hours, days, or even weeks apart. This is the major Latency of reporting; the fact that by the time a report is generated and reviewed, the conditions that caused information to appear on the report may have occurred so long ago as to make it impossible (or at least very difficult) for an organization to intelligently act on and respond to that condition.

Critical information must be detected and acted on as soon as possible.

Some organizations try to address this Latency by running reports very frequently, such as every 10 or 20 minutes. But who can afford the impact on system processing of running a report so often? And (even more importantly), who wants to explain to a Sales Manager that they’ve received 24 blank reports today ***just in case*** there was some recent business activity that required their attention?

Today’s departmental Managers and Executives don’t want to receive ***more*** reports, they want to receive ***more meaningful*** reports. That’s where “Exception Management Reporting” comes into play.

Exception Management Reporting is an approach whereby the conditions that warrant reporting on are watched for on a frequent basis. When those conditions do exist – and only then – does a corresponding report containing the relevant information get generated and delivered. This approach combines the best of Exception Management (see Latency #2) and Redundant Task Execution (see Latency #1).

Not only does this approach dramatically reduce the Latency between the appearance of and response to critical business data, it also frees an organization’s “Human” resources for more important tasks.

Reducing the Latency in Generating & Distributing Reports:

The best way to reduce the Latency in Generating and Distributing Reports is through Automation, but there are many different types of Automation available. Thus the first step is to identify which Latencies you wish to address. You can reduce Latencies that result from:

- Manually-submitted reports
- Distributing multiple copies of the same report (impacting server performance)
- Distributing reports without data
- Distributing reports that are not timely
- Triggering reports based on database conditions

You'll notice that the preceding Latencies are divided into three groups; each group can be addressed by different types of Automation. (Some Automation solutions address all three types of Latencies.)

Implementing a solution to begin addressing the Latencies of manually-submitted reports is fairly simple, as all that is initially required is a Report Distribution system. Such systems allow reports to be scheduled for generation and distribution on a recurring schedule, such as daily or weekly. Vendors that supply reporting software often offer add-on modules that automate the scheduling and distribution of their reports.

The ability to reduce the Latency connected with producing and distributing multiple copies of the same report is a bit more challenging. Most people reason that if 25 people need to receive the same report, the report needs to be generated and distributed 25 times. That does not necessarily have to be the case.

There are two very efficient ways to distribute the same report to multiple people; via the Internet (web browser), and via FTP (File Transfer Protocol). Some Report Distribution Automation systems have the ability to post (or "push") reports to a website. Thus if you have 25 people who need to receive the same report, you can have the Automation solution post one copy of that report to a centralized website, and direct all 25 people to that site to view (and optionally download) the corresponding report.

FTP'ing can also be used in a similar fashion. Automation can FTP one copy of the report to a centralized server, where the appropriate recipients could download it. Another option is to have Automation FTP separate copies of the report to each individual recipient's preferred location. Regardless of whether the Internet or FTP is used, email can also be used to notify the report recipients that their report is available – and even provide a URL link directing the recipient where to go to access the corresponding report.

Among the virtues of distributing reports via the Internet or FTP is that they provide the ability for the report recipients to view the report on-line. In many cases, the need for a recipient to actually print out the report is reduced or eliminated altogether.

The last area of Reporting Latency is the most difficult to achieve, but at the same time it also provides the greatest value, as it dramatically improves the quality of the data being reported on. The Latency in this case stems from the inability of an organization to generate reports when they are needed, and only on the information that requires attention. Called "Triggered Reporting" or "Automatic Exception Reporting", this is the ability to base the generation and distribution of reports on the presence of certain conditions of business data.

Automation that provides this ability eliminates the problem of blank reports, as reports get generated only if the appropriate conditions exist. Automation that provides this ability also eliminates most of the lag-time between the presence of “reportable data” and the generation and distribution of the corresponding reports.

And, Automation that provides this ability can be restricted to producing reports based on very specific data (such as an inventory item within 10% of its re-order level), or on very specific thresholds (such as a support rep with more than 6 high priority calls).

But keep in mind that Automation that supports scenarios like these is only as valuable as an organization’s ability to identify what those scenarios are.

When considering the use of Triggered Reports, it is essential that an organization first consider which reports would benefit from being more time-sensitive, and then consider what business conditions should trigger a corresponding report. Once an organization has identified those items, the return on investment from a Triggered Reporting Automation solution is easy to justify.

Latency #6: Integrated Data Analysis & Response

Definition: The need to automatically and dynamically analyze data between multiple business applications and (if appropriate) trigger the appropriate automated workflow.

There’s a reason why this Latency was saved for last: it’s most difficult to achieve. Most organizations have not considered “integrated data analysis & response” in terms of its embedded Latencies because everyone assumes that those Latencies are a given.

This assumption is based on the fact that most businesses use multiple software applications and the very existence of that variety of software demands that certain Latencies cannot be avoided.

The truth is, they can.

Examples:

- ➔ An organization has a “Sales” application and an “Accounting” application. They do not share data. There is no way to know when a client’s pending sales (in the Sales DB) would put them over their credit limit (in the Accounting DB).
- ➔ An organization has 3 databases in which employees record time spent on different activities. There is no way to automatically total the times from all 3 databases per employee and see if that employee is above or below certain thresholds.
- ➔ If someone in the Finance department puts a client on ‘Credit Hold’ (in their Financial DB), there is no way to ensure that the client’s salesperson knows about it.
- ➔ It’s impossible to spot interdepartmental “trouble spots” e.g., customers who have ordered less than \$5k of products but have called into Support more than 25 times over the last year.

Business Impacts:

If it's difficult for an organization to keep on top of the activities of just a single business unit or department (such as Sales, Finance, or Customer Support), think of the challenge of trying to maintain a high level of awareness of the "convergence points" of business data between those departments.

From a "lag-time" perspective, consider this:

If the average department has the time and energy to analyze and respond to their business data once a month, how often will an organization as a whole have the time and energy to analyze the combined business data between their departments? Quarterly? Yearly? Ever?

And lastly, who within an organization has the time, energy, or expertise to analyze cross-departmental business issues? The CEO? CFO? Anyone?

Any organization that focuses all of their attention on departmental goals, exceptions, and thresholds is being short-sighted. An individual client may appear to be a spectacular customer from the perspective of Sales, but might be a nightmare to Customer Service and Finance. The client might be a late-payer, always near their credit limit, and call into Support ten times a day.

A complete picture of this client cannot be gleaned from any one department. The truth to that client's status and value lies in the sum of all their activities.

The impact of an organization's inability to see this can be significant.

Consider a customer whose Sales profile is rosy but whose Financial profile is bleak. Without cross-departmental awareness, Sales staff can spend extensive time and effort closing a deal only to have it put on hold by the members of Finance. Lost time? Wasted Effort? You bet.

The problem with most cross-departmental situations like this one is that they typically do not reveal themselves until it's too late. It's only after Sales has worked hard to win an opportunity and thrown it over the wall to Finance that a problem is detected. And although the sale itself may be halted, the time spent by the Sales Team is gone, never to be regained. That's a Latency that no company can afford on a frequent basis.

More and more companies are understanding that cross-departmental analysis is an essential part of managing their business. However, the cross-departmental analysis methods adopted by some companies often result in even greater Latencies than before. The problem here is that many cross-departmental analysis tools themselves cause more Latency than they relieve.

The reason for this is that most organizations assume that they need two, three, or more Automation solutions to adequately address the need for cross-departmental analysis and response. These solution sets typically include:

-
1. An Integrated Data Mining tool (to combine and analyze data between departments)
 2. An Integration tool (to move the data from one application to another)
 3. A Workflow tool (to trigger the creation of follow-up actions)

Unfortunately, it is often the case that three such tools as these do not easily integrate with **each other**. And so, the time and effort required to integrate these solutions with each other and then apply them to cross-departmental awareness often results in a net loss of time rather than a net gain.

Reducing the Latency in Cross-Department Analysis

Although the challenge of cross-departmental awareness is significant, the method to address it is relatively simple and consists of two steps. The first step is to identify the departmental “convergence points” of business data.

To identify convergence points, you need to ask such questions as:

- What actions by department ‘a’ require a response from department ‘b’?
- What actions by department ‘a’ can be impacted by actions from department ‘b’?
- What thresholds which are determined by department ‘a’ can impact actions or decisions by department ‘b’?
- What actions by department ‘a’ need to inform members of department ‘b’?

Process flowcharts are an invaluable aid to identifying business process convergence points. In most flowcharts, you will readily be able to see when a decision point has been reached, and what departmental data is required to reach that decision.

Keep in mind that the **earlier** in a process that decisions can be reached, the less impact a Latency will have. For example, a process that waits for a sales opportunity to be worked on and closed before checking a client’s financial status has a much larger Latency than a process that checks a client’s financial status as soon as an opportunity is created.

The second step in addressing the challenge of cross-departmental awareness is the selection of a **single** Automation solution that encompasses all three capabilities mentioned previously.

Effective cross-departmental analysis and response requires a solution that offers a combination of Integrated Data Mining, Data Integration, and Workflow capabilities. Historically, those three functions have been available only as separate Automation solutions, but the last few years (from 2000 onwards) has seen the introduction of such consolidated solutions onto the market.

When reviewing potential Automation solutions, an organization needs to first evaluate what extent of functionality they need in each of the three areas mentioned previously (Integrated Data Mining, Data Integration, and Workflow) as different organizations require different levels of functionality in each of these areas. Once that analysis is complete, an organization can effectively choose a consolidated Automation solution that offers the precise level of functionality they require.

Summary:

When it comes to improving an organization's efficiency, today's companies have come a long way over the past 20 years. For the most part, however, the increase in efficiency has been due to Enabling Technologies that have allowed us to process information faster, store it more efficiently, and retrieve it more intelligently.

But there comes a point when companies that have already invested in Technology can gain only so much more efficiency from improvements to that Technology. And it is at that point that companies need to look inward, at their own processes and procedures, to see where delays, bottlenecks, and simple lack of awareness play major roles in impacting their efficiency.

Delays, or "Latencies" exist in all organizations, and can be readily identified in a variety of areas of a company's business. These Latencies impact an organization's efficiency in more than just the mere time they take up; they have a domino effect insofar as they are indicative of business problems that cause organizations to lose productivity, revenue, and customer satisfaction.

Latencies will never completely disappear, but then again organizations are not so much racing against the clock as they are racing against their competition. And for an organization that can adopt procedures and supporting Automation that reduces its Latencies – at least to the point of becoming more efficient and more responsive than its competition – the future is a bright one.