



INFORMATION IN REAL-TIME

RISKS AND REWARDS IN BUSINESS INTELLIGENCE

© Dan Linstedt, 2010 – This document introduces the ideas around real-time information processing in the business intelligence arena. Insights, rewards and risks are covered at a business level; including costs of acquisition and probabilities of success. If you're interested in RT for BI, then this report might be for you. Source: <http://danLinstedt.com> - Coaching

INFORMATION IN REAL-TIME

RISKS AND REWARDS IN BUSINESS INTELLIGENCE

Real-time information is a necessity in today's world, and with the advent of ever improving hardware it becomes easier and more accessible to achieve. However don't let the ease of acquiring the information throw your business for a loop. There are many different difficulties that can quickly disrupt your efforts. Some of the top issues surrounding real-time information include:

- Security
- Privacy
- Cost
- Time
- Systems Uptime
- Training / Education
- Investment in Infrastructure
- Information Quality

Business Justification Required

In order to properly meet the needs of a real-time system, the business **must** outline the questions that will be answered through the information that will be made available. The cost of implementing real-time information flows can be prohibitive and can easily cause BI projects to be seen as a failure; either undeliverable or beyond scope.

Without proper business justification, a real-time information project is just another IT hobby rather than a business value-add asset.

Business justification comes in many forms, however the following requirements are necessary in order to make the Real-Time BI a reality.

- 1) The top 5 questions, problems or issues that real-time data **would solve**, or **assist** in solving.
- 2) The business must identify a sponsor, someone with enough budgetary authority to approve and review IT recommendations, and business changes necessary to get the right information to the right place at the right time.

RIGHT-TIME/REAL-TIME

In these times, information availability is critical.

However what is more critical is the ability for business and IT to work together to successfully implement a plan that delivers real-time information in the right-time.

This article introduces the business to the process of ensuring right-time information delivery in BI.

- 3) The business must clearly define what **real-time** means, in other words – is it a 2 second response time for answering one question, while it's a 5 minute response time for another? Or is it current information flow, direct from source to display? The definition may be different for each question on the list!
- 4) The business **must** set budgetary limits for the project **before** the project begins.
- 5) The business **must** outline how **success** of the real-time project will be determined. These are the KPA's (Key Process Areas) and KPI's (Key Process Indicators, Metrics) by which IT and the project will be guided.

Assessing Feasibility

IT leads and business users must work together to assess the feasibility. One of the biggest mistakes that real-time projects make up-front is defining goals and objectives but not assessing them to see if they are realistic. The feasibility of meeting the KPI's for each question must be examined. In other words, if the business determines that question #3 must have an instant response time (meaning an event happens, and the business wants their BI solution to see it immediately), but there is no infrastructure to deliver the data in the required time frame (milliseconds), then that part of the project runs a high risk of failure to meet the requirements. **UNLESS** the business has pledged to invest the appropriate monetary resources, time, training, and technology to make it happen.

The feasibility of each question, along with the **risk of not building the solution** must be examined, measured and weighed before the project starts. If the information that is being delivered in real-time means that you can save a person's life, it's quite possible that the business can justify the massive inflow of funds needed to build the appropriate people, processes, and technology. However if the question is determined to have a low impact ROI for the organization, then maybe the question needs to be removed from the requirements for the project.

THE REALITY OF REAL-TIME INFORMATION

Not many business users understand that real-time costs money. Furthermore, the closer the information is to a zero latency deliverable, the **higher the cost** on an **exponential** scale! The same can be said for **difficulty to implement**, complexity of maintenance, challenge of integration. They all increase exponentially in challenges to IT as the latency reaches near zero. In real example of cost, suppose it costs \$125k to build a system that delivers data every 5 minutes. To change that system to deliver data twice as fast (every 2.5 minutes) will cost exponentially more. \$125k to the power of 1.1, perhaps to the power of 1.5 or more; the exponent will change with the cost of the technology, training, and knowledge needed to reach the goal.

To take it from 2.5 minutes, down to 1.25 minutes in latency will go up again on the exponential scale. This is probably the most important factor to realize. **Near zero latency (Real-Time) information systems are expensive.** There is simply no way around it. If you can live with 5 second refreshes, and 5 second refreshes are considered affordable by your corporation, then don't strive to deliver 4.5 or 3.5 second refreshes.

PART AND PARCEL – THERE IS NO HALF-WAY POINT

Once a business has decided to build a real-time system, there is no turning back. Why? Because it is costly to engage in, and secondly once engaged – can incur large amounts of sunk costs. You can build a real-time information flow to answer **specific** questions. The project cannot afford to be nor

become a **generally applied** project. It is not your typical BI system where you can have 100 to 1000 users asking different questions and getting their answer sets. Real-time systems are **tactical** and as such are point-focused to deliver **specific** answers.

If the business commits to a real-time information system (data warehouse, business intelligence, alerts, etc...), then the business must also realize they are committing to: training the people, hiring the right experts for guidance, corrections to operational systems (to provide the right data at the right time), and infrastructure re-tooling. Often times the business conveniently forgets that their IT staff have been trained to handle batch data, and downtime is OK. These IT resources need re-training! In a near-zero latency system, they need to understand that being on-call is important, that zero downtime is the only way to maintain the system, and that batch data handling architectures **don't work** in a real-time situation. The business **must** commit to re-training their IT staff to handle real-time information flows, alongside of committing monetary resources for real-time data feeds and infrastructure needs.

There is **no half-way point** of committing to a real-time information system. It's **do it, or don't do it**. It is a yes-no project. What works in "batch" and "strategic" information systems **does not work** in real-time systems.

BUSINESSES TAKE HEED OF DATA AVAILABILITY

Another common mistake of business users is to choose to ignore data availability. You **cannot** ignore data availability. If your source system delivers messages every 5 seconds, then you **cannot** have a real-time information flow that updates your mobile systems every 2 seconds! Studying the data availability needs to happen during the feasibility of answering the real-time information feed questions (project justification). You (as a business user) have three choices: a) invest money to upgrade your source system and infrastructure, b) downgrade your requirement to meet the data availability or c) cancel that part of the project.

ASSESSING DATA PRIVACY

Data travels faster and faster over the wire, and increasingly, wireless. Each data point (generation, transfer, reception) is a cause for concern in the security and privacy areas. Increasing the rate at which data flows only increases the risk of a security breach. What's worse, is that it increases the risk of a security breach going un-noticed or un-detected. The faster information travels, the more stringent the requirements should be to keep data safe and private.

Imagine this situation: alerts are due to be refreshed every 2 seconds. The data is prepared, the list of devices is setup for action and the system is turned on. Two seconds go by, the first alert is delivered – everyone celebrates. Five seconds goes by, and someone discovers an unauthorized device on the alert list; **too late... the data has been distributed... twice!** There is damage control to consider now.

Part of the feasibility study needs to address the "what to do to test", how the testing cycle and release cycle will work, and how the information will be governed for distribution. The business will need a fast reactionary cycle, along with a highly defined chain of command for addressing real-time privacy/security breaches. This puts everyone **on-call**, including the business, not just IT. The business users **must be available** for escalation when called upon to deal with these situations. Note that this too, increases the **cost** of the real-time systems.

Information Quality is Inversely Proportional to Real-Time Latency

Business users **TAKE NOTE!** The faster the data moves through the systems, the less time IT and Business users have to clean it up! This potentially means **more dirty data quicker**. This can be bad for business (just think of a rogue employee, or a system that has gone haywire). There are rules, transformations, and quality checks that can happen in 5 minute refreshes that will **never complete in time** for 2 minute refreshes.

When a system is delivering up-to-the-second information, it often is **not capable** of processing that information through most functions (quality or otherwise) before delivery. Transactions generated are the transactions you receive in a real-time information flow. IT cannot take responsibility (in most cases) for the quality of the information flowing through a real-time system. It is up to the business to define **what real-time means**, along with the acceptable thresholds (frequency of delivery). It is also up to the business to determine the acceptability thresholds for the quality of the data that they can tolerate.

As a business user, it is up to you to understand that near-zero latency means: what you put in, is what you get out... Now, take note, if you can live with 2 second time delays there are some systems that will work with the information and do it's best to clean it up along the way. Most systems below two to five second refreshes cost hundreds of millions of dollars to implement properly.

REAL-TIME INFORMATION FLOW IS NOT YOUR TYPICAL ETL CYCLE

When you deal with real-time information, you are dealing with individual transactions rather than batches of transactions. There is no **time** to stop and smell the roses (as it were), to run multiple aggregations, data quality, cleansing routines, or even potentially some lookups. The information movement needs shift. Most IT departments now require an understanding of BPEL (business process execution language), and the use of a tool that is in the BPM category. Something that flows and routes transactions around a publish/subscribe message bus rather than an ETL cycle.

Don't misunderstand me, ETL is still important – as you still have batch data and other integration problems to solve. Now however IT must add to it's tool set a business process queueing environment. There are some tools in this category like MQ Series, Tibco, Vitria, and WebMethods that provide transaction routing. This is where the business logic must take place to **change or alter** the data. The terminology used in this case is called **routing**.

Some traditional ETL tools have these capabilities through extensions, but, as mentioned earlier – the IT staff need to be re-trained to deal with the real time data. Traditional batch techniques often **do not apply** to processing real-time information feeds.

AFTER ALL THIS, IS THERE HOPE?

Yes, take heart, there is hope. It simply is a matter of adjusting your expectations, changing your terminology from real-time and instead, talking about **right-time**. This is the appropriate way to make your BI information flow work for you. Most systems can operate at 3 to 5 second intervals without too much pain, and a minimal of cost increases. It's when up-to-the-second transactions are necessary that the cost increases dramatically. **Right-Time Information Systems** can be built successfully, if the guidelines outlined in this article are followed.

LESSONS LEARNED ON THE WAY TO RIGHT-TIME BI

Here is a summary of some of the lessons learned along the way from a number of industry veterans:

- Don't shortcut the project, once established, studied, approved and funded, **do not** attempt to shorten the delivery cycle by "adding more resources" to the staff
- Establish proper goals, and objectives – and put in place the proper metrics to measure success along the way
- Make sure the IT staff (and a few technical business leads) are properly trained before the project begins. It will save everyone time, money, and headaches down the road.
- Architect the systems solution first, before any implementation design is chosen.
- Decide on the right-time or acceptable level of delivery, ask the question: What can we live with? Don't ask the question: What is it I think I need?
- Accept that data may arrive dirty or inaccurate, and on top of that, accept the realization that you might have to increase the data quality of your transactional systems **in order** to receive the benefits at the delivery end.
- Build appropriate escalation paths and procedures, **and don't stop at identifying IT resources only!** Include your business resources as a part of the success plan
- Establish proper risk assessment and mitigation strategies; analyze the risk of the project at frequent intervals. Get involved with helping the IT team meet the objectives.
- Understand that faster data costs more money, and that the investment includes **people, processes, training, and technology**. That it's not simply an investment in "yet another IT project" that this one is different – and **requires** the **active** participation of business users.

Finally, focus on answering one question at a time and then measuring it's success based on pre-established KPA's and KPI's. Remember that answering the first question is always the hardest – because that's where the investment in training, infrastructure, preparedness, and architecture happen. It's an investment in people, process, and time. Generally once the first question is answered, the second one becomes just a bit easier (as long as the latency is the same or lower).

I provide a full range of consulting in right-time information processing systems. I've spent hundreds of hours helping the US Department of Defense, US Navy, US Army, and a variety of other organizations with systems that reach thousands of end-users. Please feel free to forward comments and suggestions to me, and if you're so inclined – I'd be happy to chat with you. Simply pick up the phone and call.

Daniel Linstedt
President, Empowered Holdings, LLC
Web: <http://empoweredHoldings.com>
Tel: +1 802-524-8566