

# InfoSphere Streams Is A Game Changer

Merv Adrian, Principal, IT Market Strategy  
[www.itmarketstrategy.com](http://www.itmarketstrategy.com)

IBM has made it clear that [InfoSphere Streams](#), the commercialized part of the System S research project that has been underway for some years, is a priority, and they are committing substantial investments to it. In fact, the release was hurried a bit, as I noted in my (hopefully) [humorous post](#) about naming for complex event processing (CEP) and related technologies. At a major financial analyst meeting in May 2009, CEO Sam Palmisano called it out as an IBM opportunity, and Software Group honcho Steve Mills listed it as one of four themes within his topline Information Agenda message. That kind of push makes things happen.

I've seen skepticism about this from people I respect like Curt Monash ([here](#) and elsewhere), Neil Raden ([here](#)), and Peter Thomas ([here](#)), and it's not surprising - there is much evangelism to be done here to establish the vision. What are the use cases? IBM is starting to describe them based on real, money-making engagements, and they will continue to. Let's look at a few things that will drive this technology into wider use.

- **The world is being instrumented.** I hardly need to give you the numbers - they're everywhere. It's not just radio telescopes (although that's cool) - it's traffic cams. RFID tags. Trades on financial exchanges. Transactions on electronic consumer systems. And that's only the beginning. There is a class of information for which the right answer is not "let's put it into a box and then look at it." Michael Stonebreaker, a seminal thinker and innovator in database, said as much when he started [StreamBase](#) a few years ago, and that firm's lack of breakthrough success so far only indicates to me that they were a bit ahead of the market need. We are deploying more instruments streaming data all the time, and that will only accelerate.
- **Early use cases have proved out.** Admittedly, some are still rocket science, but in the financial industry (where StreamBase is doing most of its business) the notion that one might usefully filter and react to things as they go by, without building a place to keep them first - or at all - makes sense. And in everyday monitoring-based use, who cares about last month's traffic patterns, or the shipping record for that package last month? Somebody does - but not for every application we'll associate with that business activity. And certainly not for all the interesting ones.
- **Tools are emerging.** There are SQL-based approaches, including the [work being done at SQLStream with standards](#), and IBM has introduced an Eclipse-based IDE, with its language called Stream Processing Application Declarative Engine (SPADE), that includes visual metaphors that are lightyears away from the thorny approaches of the earliest uses. And with rules engines and decision automation climbing executive wish lists, the combinations could

be compelling. Object-oriented databases were a zero-billion dollar market and always will be, because only scientists used them, and they weren't moving toward mainstream use. That is simply not the case here.

- **Complementary technologies are available.** IBM is coupling InfoSphere Streams with their recently acquired [SolidDB](#) in-memory database for some classes of applications. They and other vendors have fashioned adapters to simplify connections to existing industry-standard streams of information in finance, logistics and healthcare, again, shortening the time to deliver something meaningful. And more streaming standards will emerge.
- **Expertise will drive proliferation.** Complex software is everywhere, and it creates employment opportunities for SIs. Who has one of the biggest such organizations? IBM. Who has demonstrated world-class capability to replicate hardened IP based on engagements they've done a few times? Ditto. As the flywheel picks up speed, watch out.

So, if you think I'm succumbing to the hype, so be it. I believe that we are at the beginning of a new class of applications that will be as fundamentally different from database-based ones as those were from the kind that made single passes of files in the early days of punch-card systems. As each new level of abstraction emerges, we think differently about the world, and the art of the possible. If a baseball player had to compute angles and speeds and gravity and wind and the weight of his bat before swinging, how many hits would he get? And yet, without storing data and doing queries against it, he responds to events with flashes of partial data, context, and some rules, and knocks it out of the park sometimes. Oh, and misses sometimes too. (Raden, that sports analogy was for you.) But that just means we factor that into our thinking about what we try to accomplish, and what's at stake.

Stream-based computing will change the world, and the power of IBM's engine will be a big accelerator of the change. Technology, marketing, sales, services add up. Steve Mills' organization provided 43% of IBM's profit in 2008, and his belief that, as he has told me, "our investments position IBM to do things nobody else can do" is spot on. Mills told the financial analysts at the May 13 meeting that he is focused on "better optimization of the business decision process and becoming more predictive." Elaborating on that, he spoke about federation, pattern recognition and prediction - with streams as a key enabler. This is visionary, and in sync with the "Smarter Planet" theme, so it will be relentlessly sold at the highest levels of the world business community. IBM changes markets. What's come so far has been prologue. Change is gonna come.

Published as <http://mervadrian.wordpress.com/2009/05/22/infosphere-streams-is-a-game-changer/> on May 22, 2009.