CITO Research | March 2014

The Hidden Real Time Computing Foundation That is Disrupting the World

Share:

By now, it's evident that the speed of business is heading inexorably toward "real time," and that means that industry by industry, a specialized infrastructure is emerging to support disruptive forms of business that has the net effect of driving higher yields for an entire ecosystem. Most of this happens quietly in the background, but without latency-reducing direct interconnections, the recent transformations of financial trading and online advertising would have been impossible. Since many more industries will follow finance and advertising, it is imperative for leaders to understand how to make an industry operate in real time.

The push to real time started with finance. Since electronic trading became the standard, buyers and sellers began looking for ways to accelerate transactions. At first, firms had no choice but to purchase expensive dedicated lines from a data center provider in order to exchange, and often point-to-point connections between firms. With the advent of carrier-neutral interconnection points such as Equinix, many firms ultimately decided to co-locate their trading engines next to each other in secure, neutral, high-performance data centers. With ubiquitous, low-cost computing in the cloud, the need for completely dedicated infrastructure has diminished. But the advantage of proximity has only increased, and the same speed requirement is emerging for other businesses, such as advertising.

We don't know when and which industry will develop instantaneous connectivity needs, but when it happens, it will seem like it "happened overnight," and no one in charge of technology will want to be on the wrong side of the transition.

The Evolution of Ad Exchanges

Ad exchanges emerged because of inefficiencies in the previous bilateral model.

Previously, an online outlet would have a given number of open slots for advertising, and would sell those slots directly to one advertiser. But, as the online ad market grew, other advertisers now wanted those same slots, as well as slots on other Web properties. Soon, because of the sheer scale of advertising inventory, ad networks arose to manage these transactions. Eventually, multiple ad networks popped up, and they aggregated into centralized, automated trading venues.

Now, extremely granular targeted ad buys are possible. In the past, advertisers would buy ads sold in increments of thousands of impressions, but there was no way of knowing that the ad was relevant to only 20 out of every 1,000 impressions. Ad exchanges made it possible to sell impressions to individual users in a specific demographic at the exact moment when they appeared on pages while searching under relevant terms, or clicking on relevant articles.

Much like the financial trading world, the presence of a viable exchange also spawned several other layers of intermediary tools that added more value for buyers and sellers. A Sell Side Platform (SSP) helps publishers sell inventory through both ad exchanges and direct selling. A Demand Side Platform (DSP) conducts bidding on behalf of advertising firms. In just a few years, this system has become just as complex, and significantly larger in scale than the financial trading industry, simply because at the end of each trade there are many more individuals involved. Everyone who clicks on an ad generates data. For context, in 2011 the combination SSP/ad exchange Rubicon Project processed one trillion trades – twice as much as Nasdaq did that year. The volume has only continued its exponential growth: as of April 2013, Rubicon Project's ad trades are six times those of Nasdaq.[1]

In this age of information, delivery speed is just as important as volume – if an ad is not delivered within a certain amount of time, it won't be displayed.

The Power of Proximity

Just as with the financial markets, the ad market became an ideal candidate for real time computing infrastructure because of the power of proximity. Placing two trading engines next to each other in an Equinix facility can bring latency down to just 2 milliseconds, compared to an average of 50 milliseconds to contact servers over public networks.

Equinix's advertising ecosystem, Ad-IX, is a neutral co-location site with carrier-grade connections between racks of equipment that can be managed and hosted by customers themselves or on behalf of customers by Equinix service partners.

[1] http://www.rubiconproject.com/press-releases/rubicon-project-drives-automated-advertising-revolution-to-next-level/

Here's an example of how Ad-IX delivers value. BrightRoll, a real-time bidding (RTB) platform for digital video advertising, situated its production systems in Equinix's interconnection platform and was able to cut its network latency by 80 percent, resulting in a bid rate improvement of 48 percent. This led to a corresponding increase in potential revenue for all parties transacting with BrightRoll. Seeing the value of this, other related service providers began to connect into Ad-IX as well.

Conclusion

Milliseconds matter in digital business. The latency-reducing power of proximity is key to enabling technology for real-time business model innovation. Any business that conducts transactions over the Internet could potentially benefit from direct interconnections with counterparties and value-added service partners. It's hard to say which industry will be the next to move to a real time ecosystem – but it's fair to say none wants to be the last.

CITO Research

CITO Research is a source of news, analysis, research and knowledge for CIOs, CTOs and other IT and business professionals. CITO Research engages in a dialogue with its audience to capture technology trends that are harvested, analyzed and communicated in a sophisticated way to help practitioners solve difficult business problems.

Visit us at http://www.citoresearch.com