

Crossing the chasm

Existing risk management information systems proved too fragmented and cumbersome to meet decision-makers' requirements in the crisis. **David Rowe** argues that a major reappraisal is required

Myron Scholes recently spoke about a concept he calls "volatility time". By this he means how quickly decisions need to be made to be effective. His point is that when volatility is high, it shortens the clock time available to make a decision. In this environment, delay can make a correct decision moot if the resulting action is taken too late to avoid losses or other damage. In effect, volatility time accelerates relative to chronological time.

A related concept is "risk management clock-speed" – an idea I first encountered in an exchange with Keith Smith, a risk consultant at Strategic Thought Group.¹ This is essentially the rate at which risk information is assembled, processed and made available to risk managers to support their decisions. Slow risk management clock-speed can be fine when volatility is low and volatility time is unfolding at a leisurely pace. Unfortunately, when volatility time accelerates, it can overwhelm a risk management information system operating at slow clock-speed. In a crisis, decisions are required in minutes or hours rather than days or weeks. In this environment, decisions are made in an increasingly heuristic fashion, informed by information that is either stale and/or poorly structured to address the issue at hand.

The problem of slow risk management clock-speed is a pervasive legacy of information systems that have often grown by accretion over years, if not decades. Global communication bandwidth was a serious constraint as recently as 10 years ago, and systems were designed to economise this comparatively scarce resource.

In addition, most source systems had no message-based means of transmitting incremental event-driven output to external systems. As a result, periodic (usually daily) production of batch files became the standard method for feeding data into central risk systems. Even in the best of circumstances, this approach meant data would be at least 16 hours out of date. For a global trading operation with a daily close at the end of the New York business day, data for the opening of trading in Asia would often be from two days earlier.

A second problem that slows risk management clock-speed is incomplete data coverage. When a

counterparty encounters serious problems, having complete information on exposure in all areas of business is crucial. We know that when Lehman Brothers failed, many firms were unable to determine the full extent of their net exposure for days or even weeks. This hampered their ability to act in a way that would limit losses. Even in institutions where aggregate exposure was known at the enterprise level, however, it was not always easy for individual trading desks to isolate the specific trades that were now in default. This slowed the process of executing the necessary hedges to rebalance the book and created unwanted market risk in an extremely volatile environment.

A final problem is poor organisation of the data and limited ability to tabulate and analyse non-standard aggregations. UBS had a significant centralised data effort that allowed it to simulate the impact of many macroeconomic events. Unfortunately, US housing prices were not among the variables that could be evaluated in this way.

The fragmented data consolidation and risk systems in place at most banks, especially the largest global players, require considerably more than marginal improvements to achieve fast risk management clock-speed. Significant improvements in timeliness will require event-driven messaging rather than periodic batch file transmission. They would also require significant revamping of their database design to allow for rapid responses to non-standard queries.

Metaphorically, there is a chasm dividing the typical existing risk information infrastructure from a system that is both up-to-the-minute and flexible enough to supply needed non-standard information in a crisis. It is obviously unrealistic to think that any institution will resolve this problem through a crash programme, nor is such an approach even advisable. What institutions can and should be doing is to recognise the magnitude of the problem and begin to establish a long-range plan to address it. This will inevitably require developing a vision for a coherent architecture that allows functional modifications and extensions without creating ever-greater maintenance problems. Central to success in such an effort will be a senior executive who can balance issues around analytical trade-offs, computational demands, requirements for timeliness and budget realities (*Risk* December 2004, page 71). What should not be acceptable is simply ignoring how far current risk management information systems fall short of what is required until the next crisis drives the point home... again. ■

David Rowe is executive vice-president for risk management at SunGard. Email: david.rowe@sungard.com

¹ Smith actually uses the term 'risk clock-speed' but I think 'risk management clock-speed' better describes the useful idea he advances