Modeling in the Cloud

By Trevor Howes

Actuarial modeling teams are under pressure from many directions. Standards and regulations are changing everywhere, imposing new methods, more granular models and assumptions that must be updated to reflect current estimates. For International Financial Reporting Standard (IFRS) 17,[1] and U.S. Generally Accepted Accounting Principles (GAAP) with targeted improvements, models must also be run again many times within a reporting cycle to analyze and decompose the causes of change for disclosure purposes.

With the pace and magnitude of change, governance of model changes and control over the process of model improvements becomes an even greater challenge. At the same time, actuaries' work is subject to increased scrutiny.

> Massively scalable cloud processing has come online that can enable faster actuarial model run completion ...

It would be wonderful if actuarial modeling teams were given more time to perform the required calculations and investigate and test the results for quality. While that is an attractive idea, it is not going to happen. Reporting windows are narrowing, more post-processing of actuarial calculations are needed, and data aggregation and disaggregation must also occur. Actuarial models are being caught in a perfect storm of change.

However, there is good and timely news from the technology side that may help address these pressures and challenges. Massively scalable cloud processing has come online that can enable faster actuarial model run completion without the commensurate capital and operating cost investment. Here is why cloud processing options can be such a benefit.

Actuarial models, especially those involving stochastic processing and hedging calculations, have required increasing computational power in recent years. The solution thus far has been to find ways to distribute the model calculations over server farms, organized as grids. This allows many independent cores to work in parallel to finish a complete run in a fraction of the time that a single core, or a single computer with multiple cores, requires to complete the same task.

It has worked well, to a degree, but has depended, at least in the beginning, on each company designing, purchasing, installing, configuring and operating its own private complex of powerful servers. The hardware purchased was carefully configured in terms of the size and speed of processors, the numbers of cores and amount of disk capacity per box, the operating software and middleware, and the network connectivity to suit not only the actuarial modeling system but also other enterprise applications demanding a share of that precious resource.

The grid capacity has had to be sufficient to meet peak needs during critical reporting windows yet also be reasonably well-occupied by other tasks between these windows. As needs have rapidly and unexpectedly increased, extra capacity has been difficult to plan, fund and bring online, impacting capital budgets and straining information technology support personnel.

The public cloud has been offering convenient alternatives for massive compute capacity for some time. The combination of newly available compute instances (including GPGPU-enabled instances) and reduced prices, plus the ability of actuarial modeling systems to exploit that new cloud capacity quickly, conveniently, and securely, is coming at an optimal time.

The greatest advantage of harnessing the cloud for large-scale modeling runs probably comes from the massive processing capacity in the cloud, which can be designed and suited to actuaries' modeling requirements. Cloud users can configure the exact number and type of compute instances with the power and associated disk space that is most effective for their needs. The virtual farm can be spun up and readied in minutes. The modeling run can commence without waiting in a queue for resources to be freed up, and without the risks of other tasks and users interfering with the performance of the run.

If multiple models need to be run for different purposes, or there are other teams or applications in the same company with the same reporting windows, they can all start their jobs and have them running simultaneously because each one is in effect requesting its own independent and private farm in the cloud.

The result is not only faster turnaround, but also greater predictability of run times, and more capacity to cope with emergency reruns of critical steps.

Assuming that cloud usage of the grid computing is only charged as it is allocated to a job, there are no wasted costs

of idle resources between runs or reporting periods, and no downside to requesting as many cores as can be managed by the modeling software—provided they can be kept busy on parallel tasks.

Furthermore, the requisition of cloud resources, provisioning of the farm, transfer of the model to and from the cloud, and release of the cloud on completion can be fully automated. This results in no extra burden on IT personnel, no upfront capital costs and planning effort, and no fighting for priority and access on the new facility. Compare that to your life sharing a fixed-infrastructure or traditional server farm. As an added bonus, the cloud allows you to quickly exploit new state-of-the-art hardware without waiting for the expiry of your current equipment lease.

So if actuaries can get prepared to harness the new capabilities and elasticity of the cloud for their modeling work, they may confidently face the aggressive challenges of IFRS 17 or U.S. GAAP targeted improvements in a timely and effective way. And maybe, that same new modeling capability will help unlock new analytics, added value and transformational change. ■

Trevor Howes, FSA, FCIA, MAAA, is a vice president and actuary at Moody's Analytics. He can be reached at *Trevor.Howes@moodys.com.*

ENDNOTE

1 International Accounting Standards Board. 2017. International Financial Reporting Standard (IFRS) 17: Insurance Contracts. *https://www.ifrs.org/issued-standards/list-of-standards/ifrs-17-insurance-contracts/.*