Partially cloudy: the benefits of hybrid deployment models

By Andrew Aziz

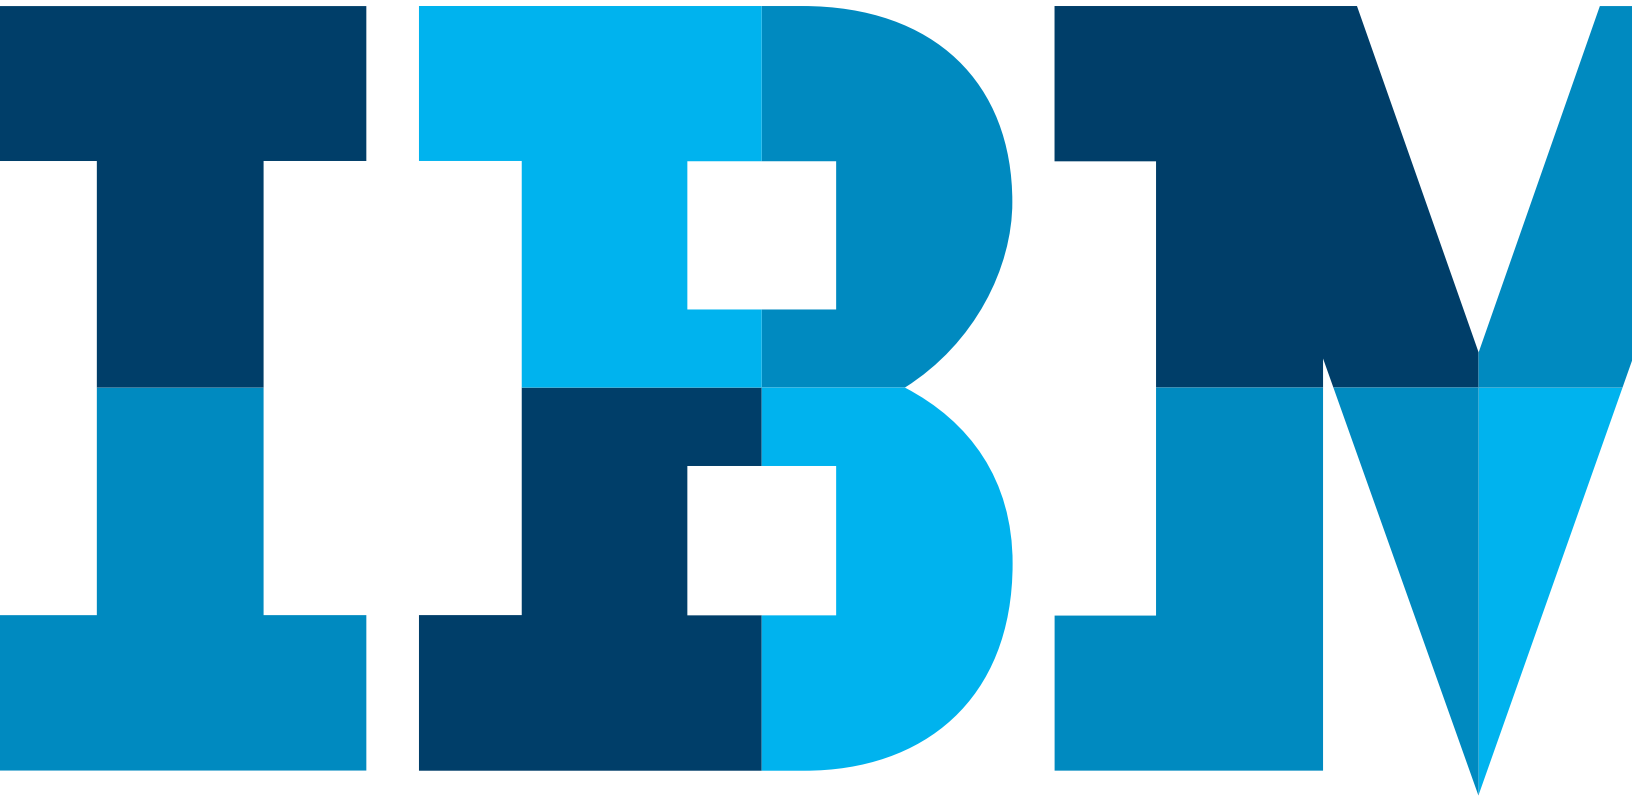

IBM

Financial institutions have traditionally faced stark choices with respect to the deployment of their risk solutions. If their desire was for configuration flexibility and transparency into the modeling and data workflow, they had been forced to bear all the potentially high costs associated with an on-premise implementation. In contrast, if their desire was to avoid the costs of data management, hardware and infrastructure support, they had been forced to accept the "Black box" implications of an on-cloud solution.

Of course, depending on the specific business requirements of individual firms, either of these extreme options may be quite appropriate. Financial institutions may sometimes prefer a full on-premise deployment. As an example, a hedge fund trading in complex structured products may have acquired a great deal of market intellectual property (IP) by trading in this asset class and, thus, desires full configuration flexibility and transparency over the risk factor and product modeling processes underlying their risk analytics. In contrast, other financial institutions may prefer a full on-cloud solution. For example, a pension plan holding generic securities or other more opaque products may be quite happy with standardized modeling choices, and look to the cloud for its inherent low cost of ownership benefits.

But the choices available are not so black and white: Today, given the continuing evolution of technology and the maturing of on-cloud capabilities, these stark extremes are no longer the only options available. Financial institutions are now able to "tap" into a much broader array of deployment options reflecting varying degrees of processing that may be allocated between on-premise and on-cloud. Equally significant, such a continuum of deployment options enables firms to benefit from hybrid deployments by allowing financial institutions to choose heterogeneous deployment options that specifically target their heterogeneous business needs.

Consider a hedge fund similar to the one described above, but trading in a broad range of asset classes; some complex, some more generic. It may be the case that, for the structured products portion of the business, the firm has developed significant IP and has differentiated itself in its unique understanding of these securities in the market. As such, the firm may wish to preserve modeling flexibility with on-premise capabilities. The same firm, however, may also hold positions in more commoditized asset classes such as equities or treasury bonds; positions requiring processing effort that the firm may find beneficial to offload onto the cloud. This example begs the question: Why bear the cost of processing generic asset classes on-premise where there is no competitive advantage to the firm in terms of specific IP? A heterogeneous deployment model allows financial institutions to better align the costs associated with preserving modeling flexibility with that portion of the portfolio — and only that portion — where benefit is truly gained from accessing this flexibility.

A similar situation, but one reflecting a different tradeoff decision, may exist if this same hedge fund also trades in an asset class that has intense processing requirements; for example, mortgage backed securities. Few would argue that mortgage backed securities represent a generic financial product and that there would be no value gained in preserving model flexibility on-premise. Nonetheless, if the processing requirement of these securities is large, the case could easily be made that the reduced cost of ownership would outweigh the value derived from configuration flexibility, leading again to a preference for on-cloud processing.

Traditionally, financial institutions that were faced with such heterogeneous requirements were either forced between the extremes of going entirely on-cloud versus going entirely on-premise or, if these options were too constraining, then being forced to manage a hodge-podge of solution silos within the firm. In the latter case, the firm would choose (or often build) on-premise solutions for some parts of the business while choosing on-cloud solutions for other parts, with no ability to aggregate consistently at the enterprise level.

However, with the functional architectures that are available today financial institutions have the means to access a continuum of deployment options, and equally important, the ability to tailor hybrid deployment models to their specific business requirements while still achieving consistent enterprise level aggregation. The key enabler for a hybrid deployment model of this nature is a functional architecture that can be componentized so that it is aligned with the key steps in the data transformation process required to produce risk analytics; the process whereby raw market data is transformed via a series of processing steps into a risk measure.

Any risk measure, whether its objective is test testing, sensitivity analysis, or a distributional statistic such as VaR or expected shortfall, is based upon a data processing workflow that involves both risk factor modeling and product modeling. At a high level, the typical sequence in this workflow can be summarized by the following series of data transformation steps:

1. **Raw Data to Derived Data**
   - Data cleansing and preparation
   - Curve construction (“derived” risk factors)
2. **Derived Data to Scenario Data**
   - Risk factor “time series” management
   - Scenario generation (risk factor projections)
3. **Scenario Data to Simulated Data**
   - Model configuration
   - Simulation (valuation over session data)
4. **Simulated Data to Risk Data**
   - Aggregation and session creation
   - Risk measures (Statistics, descriptors portfolio distributions)

Each processing step incorporates a number of critical configuration choices that are, for the most part, a blend of art and science. Many financial institutions have invested significantly in developing unique expertise around the modeling of risk factors and the valuation of individual asset classes, and require risk system capabilities that can incorporate this IP. Others are more comfortable with standardized configurations and do not require the same configuration flexibility. The value that a continuum of deployment options provides to a financial institution is the ability to invest in this flexibility within areas where it is required and not invest within areas where it is not.

A functional architecture that can be componentized to the degree of granularity captured by this data transformation workflow provides financial institutions with a continuum of on-cloud vs. on-premise deployment choices, as well as providing a number of hybrid deployment options for any individual firm. On-cloud processing can be leveraged for each step in the overall data transformation process; with individual financial institutions being able to tap in at different stages—more upstream or more downstream in the continuum—depending upon their business requirements.

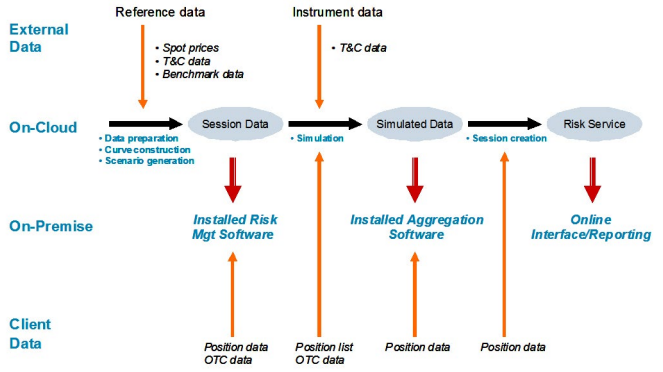

Figure 1: Flexible deployment options at each stage of the journey from raw data to risk insight.

For example, a financial institution may generally prefer an on-premise deployment for the product modeling flexibility that it provides, but does not want to bear the costs of on-premise data management, prompting the firm to offload this upstream processing step to the cloud. This firm would tap into on-cloud processing at either Step 1 or Step 2 of the data transformation continuum depending on how much control they wish to preserve in scenario

generation: Step 1 if this wish greater scenario generation flexibility, or Step 2 if they are happy with standardized scenarios (or some combination of both). In other words, the ideal deployment for this financial institution would be to install a risk software solution on-premise and subscribe to a "market data" service on-cloud, receiving derived data or derived data plus scenario data as input.

Another financial institution may generally prefer an on-cloud deployment for the lower cost of ownership benefits it provides, but is restricted from allowing portfolio holding information to leave their premises. In trying to leverage the cloud as much as possible, this firm would tap in at Step 3 of the data transformation continuum. The firm would offload all simulation processing to the cloud, while only performing aggregation processing on-premise. In other words, the ideal deployment for this financial instruction would be to install aggregation software on-premise and subscribe to a "simulation data" service on-cloud, receiving product valuations that have been simulated over scenarios as input.

Moving further downstream in the continuum, consider a third financial institution that also may prefer an on-cloud deployment for the lower cost of ownership benefits, but has no data privacy restrictions. This firm would tap in at Step 4, fully downstream in the data transformation continuum, where all processing has been offloaded to the cloud and where no software is required to be installed on premise. This financial institution would subscribe to a traditional "risk" service on-cloud, having web access to an interactive application or a series of reports. Importantly, it need not be the case that the service is a "one size fits all" model. It could just as easily be a "managed" service approach, providing more customization around the firm's needs, but still be entirely on-cloud.

Each of the previous examples describes a specific deployment model for individual financial institutions with different business requirements, with each tapping in at different steps of the data transformation continuum. However, a major benefit derived from such a componentized functional architecture is the ability for a given firm to implement a risk solution based on a hybrid deployment model, yet still achieve consistent enterprise level aggregation.

The enabler for a hybrid deployment model of this nature is that Step 3 — the simulation step in the data transformation continuum — can be comprised of multiple simulation processes running in parallel. As long as each simulation process comprising Step 3 can access a common set of scenarios as input — that is, Step 2 in the continuum — and can push results to a common destination for aggregation — Step 4 in the continuum — then consistent aggregation is achieved no matter where any individual simulation process occurs, on-premise or on-cloud.

A componentized functional architecture supports a fully decentralized network of simulation processes. This can provide key value even for those financial institutions that prefer to do everything on-premise; ranging from load balancing benefits to incorporating heterogeneous risk systems (in-house or vendor provided) in a consistent manner. Notably, the componentized architecture also provides the capability to offload some of parallel processes to the cloud in a hybrid deployment model.

A financial institution with heterogeneous business requirements can now choose which portion of the business it wishes to (or needs to) process on-premise, and offload the portions of the business it wishes to process on-cloud. This decision could be based on the "cost of ownership" vs. "configuration flexibility" tradeoff — applied across a number of dimensions including business units, geographies and asset classes — as driven by the unique requirements of the firm.

Each of the hedge fund examples described at the beginning of this article can easily be accommodated in a hybrid model by processing the treasury bonds or mortgage backed securities on-cloud, while processing their structured products on-premise.

As a final example, consider again a financial institution that may generally prefer an on-cloud deployment for the lower cost of ownership benefits it provides, but is restricted from allowing portfolio holding information to leave their premises. In this case, the firm also holds a number of Over-the-Counter (OTC) products in addition to their exchange traded securities. In trying to leverage the cloud, this firm would tap in at Step 3 of the data transformation continuum as much as possible, offloading all the simulation processing it can, while performing aggregation processing on-premise due to their privacy requirements.

The dilemma presented in this particular example is that while this deployment may apply quite well to exchange traded securities (where product description is independent from position holdings), it falls flat when applied to OTC products where the distinction between product description and position holdings is meaningless. The dilemma can be resolved, however, with a hybrid deployment whereby the OTC products are processed on-premise and the exchange traded products are simulated on-cloud; tapping in at Steps 2 and 3, respectively of the data transformation continuum.

The ideal deployment for the financial institution in this case would be to install risk software on-premise (combining simulation and aggregation capabilities) and subscribe to a data service on-cloud. The on-cloud data service would provide scenario data as input for the on-premise simulation of the OTC positions as well as simulated data corresponding to the exchange traded products. These on-cloud simulations, together with the on-premise simulations, become joint input for a common on-premise aggregation. Consistency is achieved because there is a common source of scenarios produced on-cloud in Step 2 that serve as input to two simulation processes in Step 3 (one on-cloud and one on-premise). Combined, this serves as ultimate input to a single aggregation process in Step 4 that is performed on-premise.

Financial institutions today are finding their businesses becoming increasingly heterogeneous and, in response, are seeking greater flexibility in their risk management systems and practices. While risk solutions have traditionally been compartmentalized as being delivered on-premise or on-cloud, with the ever-accelerating evolution of on-cloud technologies, a continuum of deployment options—including hybrid models—are becoming available. These options offer financial institutions unprecedented opportunities to optimize and integrate their diverse capabilities, while aligning processing costs more effectively with the areas where firms can derive the greatest value.

About IBM Business Analytics

IBM Business Analytics software delivers data-driven insights that help organizations work smarter and outperform their peers. This comprehensive portfolio includes solutions for business intelligence, predictive analytics and decision management, performance management, and risk management.

Business Analytics solutions enable companies to identify and visualize trends and patterns in areas, such as customer analytics, that can have a profound effect on business performance. They can compare scenarios, anticipate potential threats and opportunities, better plan, budget and forecast resources, balance risks against expected returns and work to meet regulatory requirements. By making analytics widely available, organizations can align tactical and strategic decision-making to achieve business goals. For further information please visit **ibm.com**/business-analytics

Request a call

To request a call or to ask a question, go to **ibm.com**/business-analytics/contactus. An IBM representative will respond to your inquiry within two business days.

IBM Corporation
Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
July 2013

Please Recycle

YTW03343-USEN-00