



# Financial Risk Management Analysis: Challenges, Changes and the Cloud

April 2016

Jeffrey Smart, Ph.D.  
Managing Director  
RIS - Market Risk Management Department  
Consumer Insurance  
+1 818 251 4513  
[Jeffrey.Smart@aig.com](mailto:Jeffrey.Smart@aig.com)

# Overview

- Our End Customer
  - Retail investors seeking retirement income for life
  
- Issuers
  - Life insurance companies
  
- Product
  - Variable annuities with living benefit rider
  - Cash flow logic described in prospectus
  - Contingent on market rates and policyholder behavior

# Business Goals for Risk Management of Variable Annuities

- Risk Management / Hedging
  - Compute values and risk statistics for retirement savings policies
  - Analytics inform daily hedge decision-making
  - Assets under management have grown from \$10 B in 2007 to \$40 B in 2015
- Valuation
  - Support GAAP financial reporting
- Capital Management
  - MRM performed analytics for CCAR (capital requirements) in 2013, 2014
  - Computationally intensive, and original driver for cloud computing
  - Capital Management analytics have moved to another department
- Support product development
  - Evaluate new products and feature changes

# Business process for hedging

- Calculate values and risk profile
  - Actual P&L: from full revaluation
  - Expected P&L: Risk sensitivities times market moves
  - Understand actual vs. expected P&L
- Identify hedge instruments (e.g., futures contracts) to offset selected market risks
- Constant demand for latest information
- Scale:
  - \$40 B Book runs overnight on 8,000 cores
  - Higher compute capacity gives flexibility

## Business process (con't)

- Increasing demands from business side
  - Product complexity grows
  - Assets under management grow
  - Reporting deadlines shrink
  - Higher model resolution / fewer approximations
- Technical constraints
  - Internal grid: too small when it's running / too big when it's idle

# Overview

- Multi-year evolution of the computational platform
  - 1: Internal data center migration
  - 2: Run a project on the cloud
  - 3: Run month-end batch on the cloud
  - 4: Run daily production on the cloud
  - 5: Decommission internal grid
  
- Considerations for internal vs. external capacity

# Step 1: Internal data center migration

- Capacity expanded
  - 10x more cores...
  - ...in 2 time zones (to cut operational risks)
  - ...for more users (quota vs. "my hardware")
- Cut month-end batch from 4 weeks to 4 days
- Multi-department grid sharply increased utilization, driving down costs

# Step 1: Internal migration (con't)

- Technical
  - New data center: infrastructure new from bare metal up
  - New capabilities: fault-tolerant, pooled resources
  - Issue resolution: infrastructure, or application?
  
- Organizational
  - End-users pooled resources; from “my hardware” to quota
  - Migration was a Company-wide initiative
  - Resources stretched thin, so stay organized
  - Successful migration is win/win for IT and business users



## Step 2: Run a project on the cloud

- Deliverable
  - CCAR Stress Tests mandated by the Federal Reserve
  - Needed 4x increase in capacity
  - High visibility project
  - Enabled critical mass of internal resources
  - Recognition for success
- Semi-annual project
  - Manual intervention feasible
- It doesn't get better (or worse) for a pilot
  - Critical, highly visible, but tweakable

## Step 2: Cloud Project (con't)

- Technical
  - Modify architecture for larger, external grid
  - More complicated than adding execute nodes to existing grid
  - Incorporate Information Security requirements
  - Modified workload (different CPU / IO balance)
  - New stresses on network, schedulers
  - Managing results (output data) is important
- Organizational
  - More people and organizations involved
  - Establish roles / responsibilities
  - New cost model
  - Internal: cost and run-time scale with number of cores
  - External: cost scales with core-hours; cut run-time at fixed cost

## Step 3: Run month-end batch on cloud

- Deliverable
  - Complete month-end runs faster
- From project to process
  - Increase reliability, stability, automation

## Step 3: Month-end batch (con't)

- Technical
  - Run internal, external grids in parallel; mitigates risk of down-time
  - Higher capacity cuts delivery time
  - Optimize ongoing process (vs. one-time project)
  
- Organizational
  - Build track record of success
  - Refine roles/responsibilities with more experience

## Step 4: Run daily batch on cloud

- Deliverable
  - Reduce cost
  - Increased capacity enables model refinement
- High demand for throughput, reliability
  - Less room for manual intervention
- Hedge on market opening –markets don't wait

## Step 4: Daily batch (con't)

- Technical
  - Focus on reliability
  - Overnight batch has less recovery time vs monthly
  - Continue to optimize for reliability, cut expenses
- Organizational
  - Continue to refine roles / responsibilities

## Step 5: Decommission internal grid

- Cost reduction
  - Measure workload in core-hours
  - Binding constraint: 8 hours for overnight runs
  
  - Internal: number of cores drives cost
  - Idle time has implicit cost
  
  - External: workload size drives cost, not core count
  - More cores / faster batch / same cost
  
- Simpler process: one environment vs two

## Step 5: Decommission internal grid (con't)

- Internal and external cost models are different

Factor	Internal Grid	External Cloud
Capacity	Fixed	Variable
Cost	Based on grid size	Based on usage
Accounting	Capital Expense + Operating Expense	Primarily Operating Expense



# Lessons Learned

- Technical
  - Make long-term plan / show incremental progress
  - Seize opportunities
  - Optimize over the right time horizon
  - Performance changes with size; rewards monitoring
  - Schedulers do not fail gracefully
  
- Organizational
  - Principles of project management apply
  - Revisit roles / responsibilities as scope grows
  - Understand constraints, cost models

Questions?