Too Big to Test

Breaking a production brokerage platform without causing financial devastation

SCMG Fall Meeting

October 20th, 2016

Kyle Parrish

Director, Technology Risk
Fidelity Investments
email: Kyle.Parrish@fmr.com

Agenda

- Who am I
- How I got into this
- The planning process
- Building the beast
- Running tests
- What we have learned
- Where we are heading
- Questions

First, a little about me…

- Active in CMG in the “Triangle” area of North Carolina
- Prior to joining Fidelity, I spent 14 years as a consultant across industries
- At Fidelity, I work in a variety of roles
  - Operational testing of brokerage platforms
  - Test Data Management
  - APM Assessment
- Have been engaged in “test in production” for approximately 5 years

A little bit about where I work - Fidelity Investments

Our systems support over…

- **24 million** individual investors
- **20.0 million** brokerage accounts
- **506,000** commissionable trades per day
- **$5.4 trillion** in total customer assets

Our technology groups support…

- 10 major businesses, all relying on shared infrastructure
- Two major data centers, each with massive mainframes
- Many different front-end channels
- Strong regulatory and compliance controls

CMG Computer Measurement Group

Disclaimer

- All of that being said…

Nothing I say or show during this presentation represents the opinions of my employer, Fidelity Investments. Everything said here results from my own individual experiences and development of best practices throughout my career in IT.

- *So sayeth me….*

The History

Where were you on the mornings of…

- August 24th, 2015?
- On May 18th, 2012?
- May 6th, 2010?
- All of September of 2008?
- Black Monday, 1987?
- October 24th, 1929??

They were not “good” days…

Dow plummets as much as 1,000 pts.

- China’s ‘Black Monday’: 2015 stock-market gains swamped by panic
- MARKET EXTRA: Exactly how panicky should we be about China?

GREATEST CRASH

DELUGE OF PANIC SELLING OVERWHELMS MARKET

19,000,000 SHARES CHANGE HANDS

PRICES TUMBLE LIKE AN AVALANCHE

D SCENES TILL GREAT INANCIERS COME TO RESCUE

The New York Stock Exchan

The worst crash in the history of Wall Street took place yesterday,

“All the News That’s Fit to Print”

The New York Times

Late Edition

VOL.CXXXVII . . . No. 47,298 NEW YORK, TUESDAY, OCTOBER 20, 1987

STOCKS PLUNGE 508 POINTS, A DROP OF 22.6%; 604 MILLION VOLUME NEARLY DOUBLES RECORD

U.S. Ships Shell Iran Installation In Gulf Reprisal

Offshore Target Termed a Base for Gunboats

By STEVEN V. ROBERTS

WASHINGTON, Oct. 19 — United States naval forces struck back at Iran today for attacks on American-registered vessels and other Persian Gulf shipping by shelling two connected offshore platforms that American officials said were a base for Iranian gunboats.

A few hours later, a naval commando detachment boarded a third platform five miles away and destroyed radar and communications equipment, Pentagon officials said.

No American casualties were reported in the actions, which occurred 120 miles east of Bahrain at about 2 P.M. (7 A.M., Eastern daylight time).

A 30-Minute Warning

American officials said the attacking force took pains to avoid killing Iranians, giving the crew on the platforms a 30-minute warning before four destroyers, stationed about three miles away, began the shelling.

At the United Nations, an Iranian delegate said “several innocent people” had been killed in the attack, but the assertion could not be confirmed.

With the bombardment, the Administration intended to send a message to Iran: The United States had shown restraint in the level of its attack this time, but might respond with greater

A Huge Blow to the Five-Year Bull Market

Dow’s Record Fall

Yesterday’s close was down 22.6 percent from Friday’s close.

10 A.M. 11 12 1 2 3 4 P.M.

1982 1983 1984 1985 1986 1987

The Dow Jones industrial average, which has been marching up since August 1982, began a dramatic fall last week that continued through yesterday when it closed at 1,738.74. Shown: Weekly close of the Dow.

Source: Knight-Ridder Tradecenter

Does 1987 Equal 1929?

By ERIC GELMAN

As stock prices soared this year, a chorus of pessimists warned that 1987 was looking more like 1929, when a stock market crash helped to usher in the Great Depression. Yesterday, after a plunge reminiscent of the worst days of 1929, one pressing question was whether the aftershocks would be as devastating to individuals and the nation.

The quick answer, many economists say, is no. The huge losses on Wall Street constitute a substantial blow to the economy at large. But there are many safeguards in place today — some instituted directly in response to the Depression — that would tend to

News Analysis

WORLDWIDE IMPACT

Frenzied Trading Raises Fears of Recession — Tape 2 Hours Late

By LAWRENCE J. De MARIA

Stock market prices plunged in a tumultuous wave of selling yesterday, giving Wall Street its worst day in history and raising fears of a recession.

The Dow Jones industrial average, considered a benchmark of the market’s health, plummeted a record 508 points, to 1,738.74, based on preliminary calculations. That 22.6 percent decline was the worst since World War I and far greater than the 12.82 percent drop on Oct. 28, 1929, that along with the next day’s 11.7 percent decline preceded the Great Depression.

Since hitting a record 2,722.42 on Aug. 25, the Dow has fallen almost 1,000 points, or 36 percent, putting the blue-chip indicator 137.5 points below the level at which it started the year. With Friday’s plunge of 108.35 points, the Dow has fallen more than 26 percent in the last two sessions.

Unprecedented Trading

Yesterday’s frenzied trading on the nation’s stock exchanges filled volume to unheard of levels. On the New York Stock Exchange, an estimated 604.3 million shares changed hands, almost double the previous record of 338.5 million shares set just last Friday.

With the tremendous volume, reports of brokers’ trades on the New York Stock Exchange were delayed by more than two hours at one point. The New

facebook

And they will not be the last…

"A recent spur in volatility across the global equity market is driving system failures at major financial institutions."

ZACKS
Our Research. Your Success.

CMG
Computer Measurement Group

The Bottom Line…

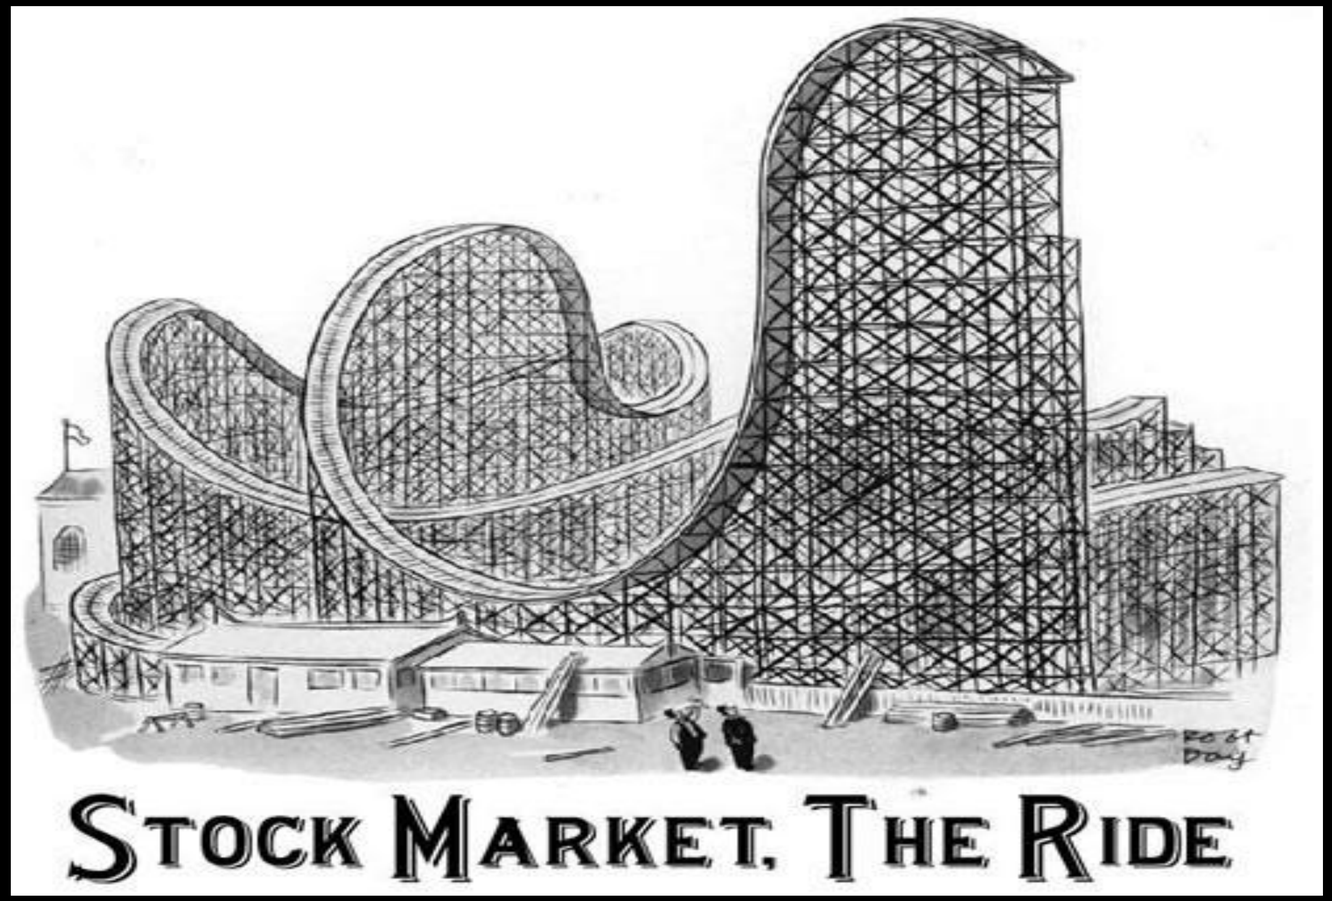

- There is no “IF” when discussing the possibility of another significant market event…only when
- The question is…

  Can we be prepared when it happens?

I believe we can
and this is what I am here to talk about

The Event

So, back to May 6th, 2010 – The Flash Crash

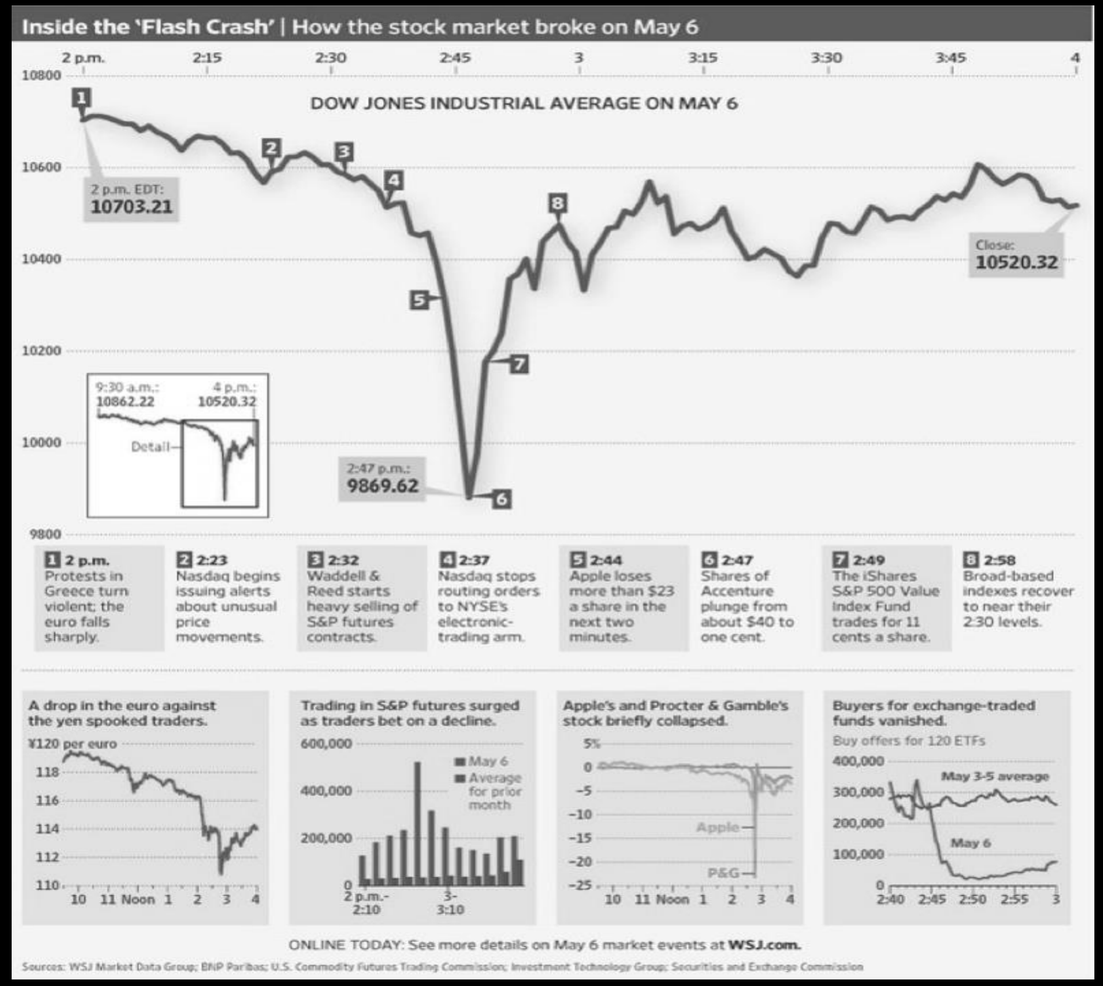

- Big events evoke big changes, and the Flash Crash of May 6, 2010 raised some eyebrows
  - *The market dropped 600 points in a span of minutes*
- Finding a new way to test was born out of this event, and has only been reinforced by events that followed
- It was a wakeup call - While technology is creating great opportunity, at the same time, it is creating great risk.

On May 7th, perf testing came to work and the management asked….

“What are YOU going to do to protect our clients?”

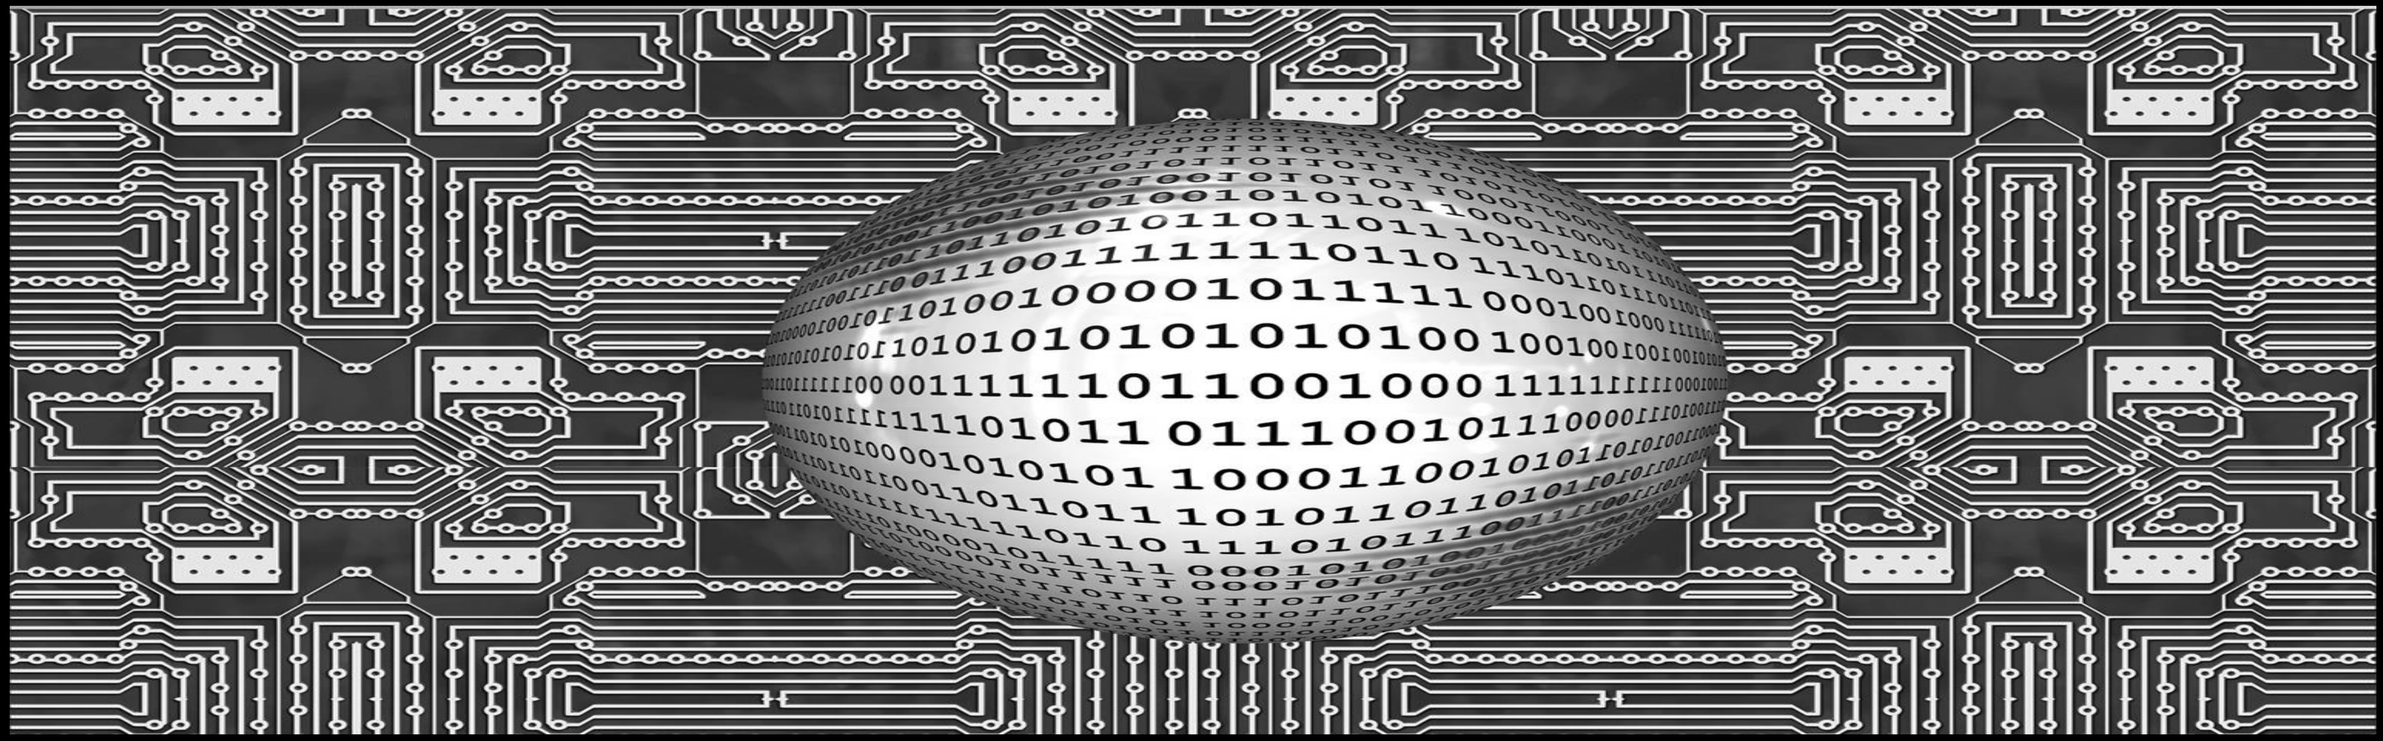

The Approach

What is the problem we are trying to solve?

- May 6th proved that…
  - Market spikes can happen at any time
  - All systems and all channels need to be flexible enough to adapt to those changing conditions
  - We could not simulate possible scenarios with the current testing protocols
  - Management wanted to pinpoint weak spots in order to mitigate the risk before any events and was, at the time, unable to identify those risks

- So, we needed to design a test protocol that…
  - Found application and infrastructure limits anywhere throughout the system
  - Enabled true end-to-end testing
  - Could repeat market open events
  - Would stress the system using many different market scenarios
  - Had the veracity of production and the flexibility of statistical models

We began thinking…

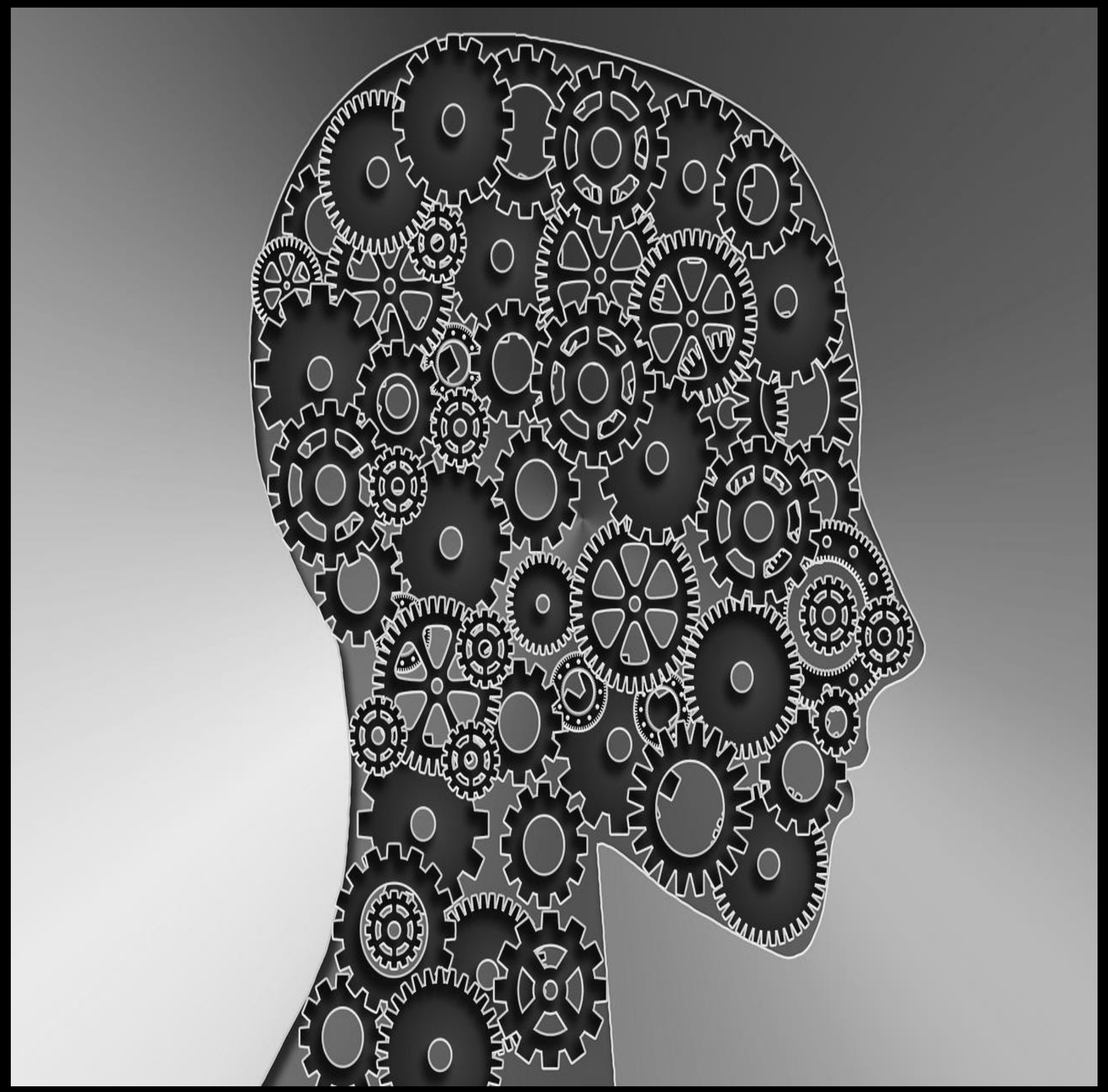

- Should we build a new environment?
- Should we work with a vendor?
- Should we repurpose an existing test environment?
- How much data do we need?
- How do we make it realistic?
- How do we test like our customers work?

And, what are we actually dealing with?

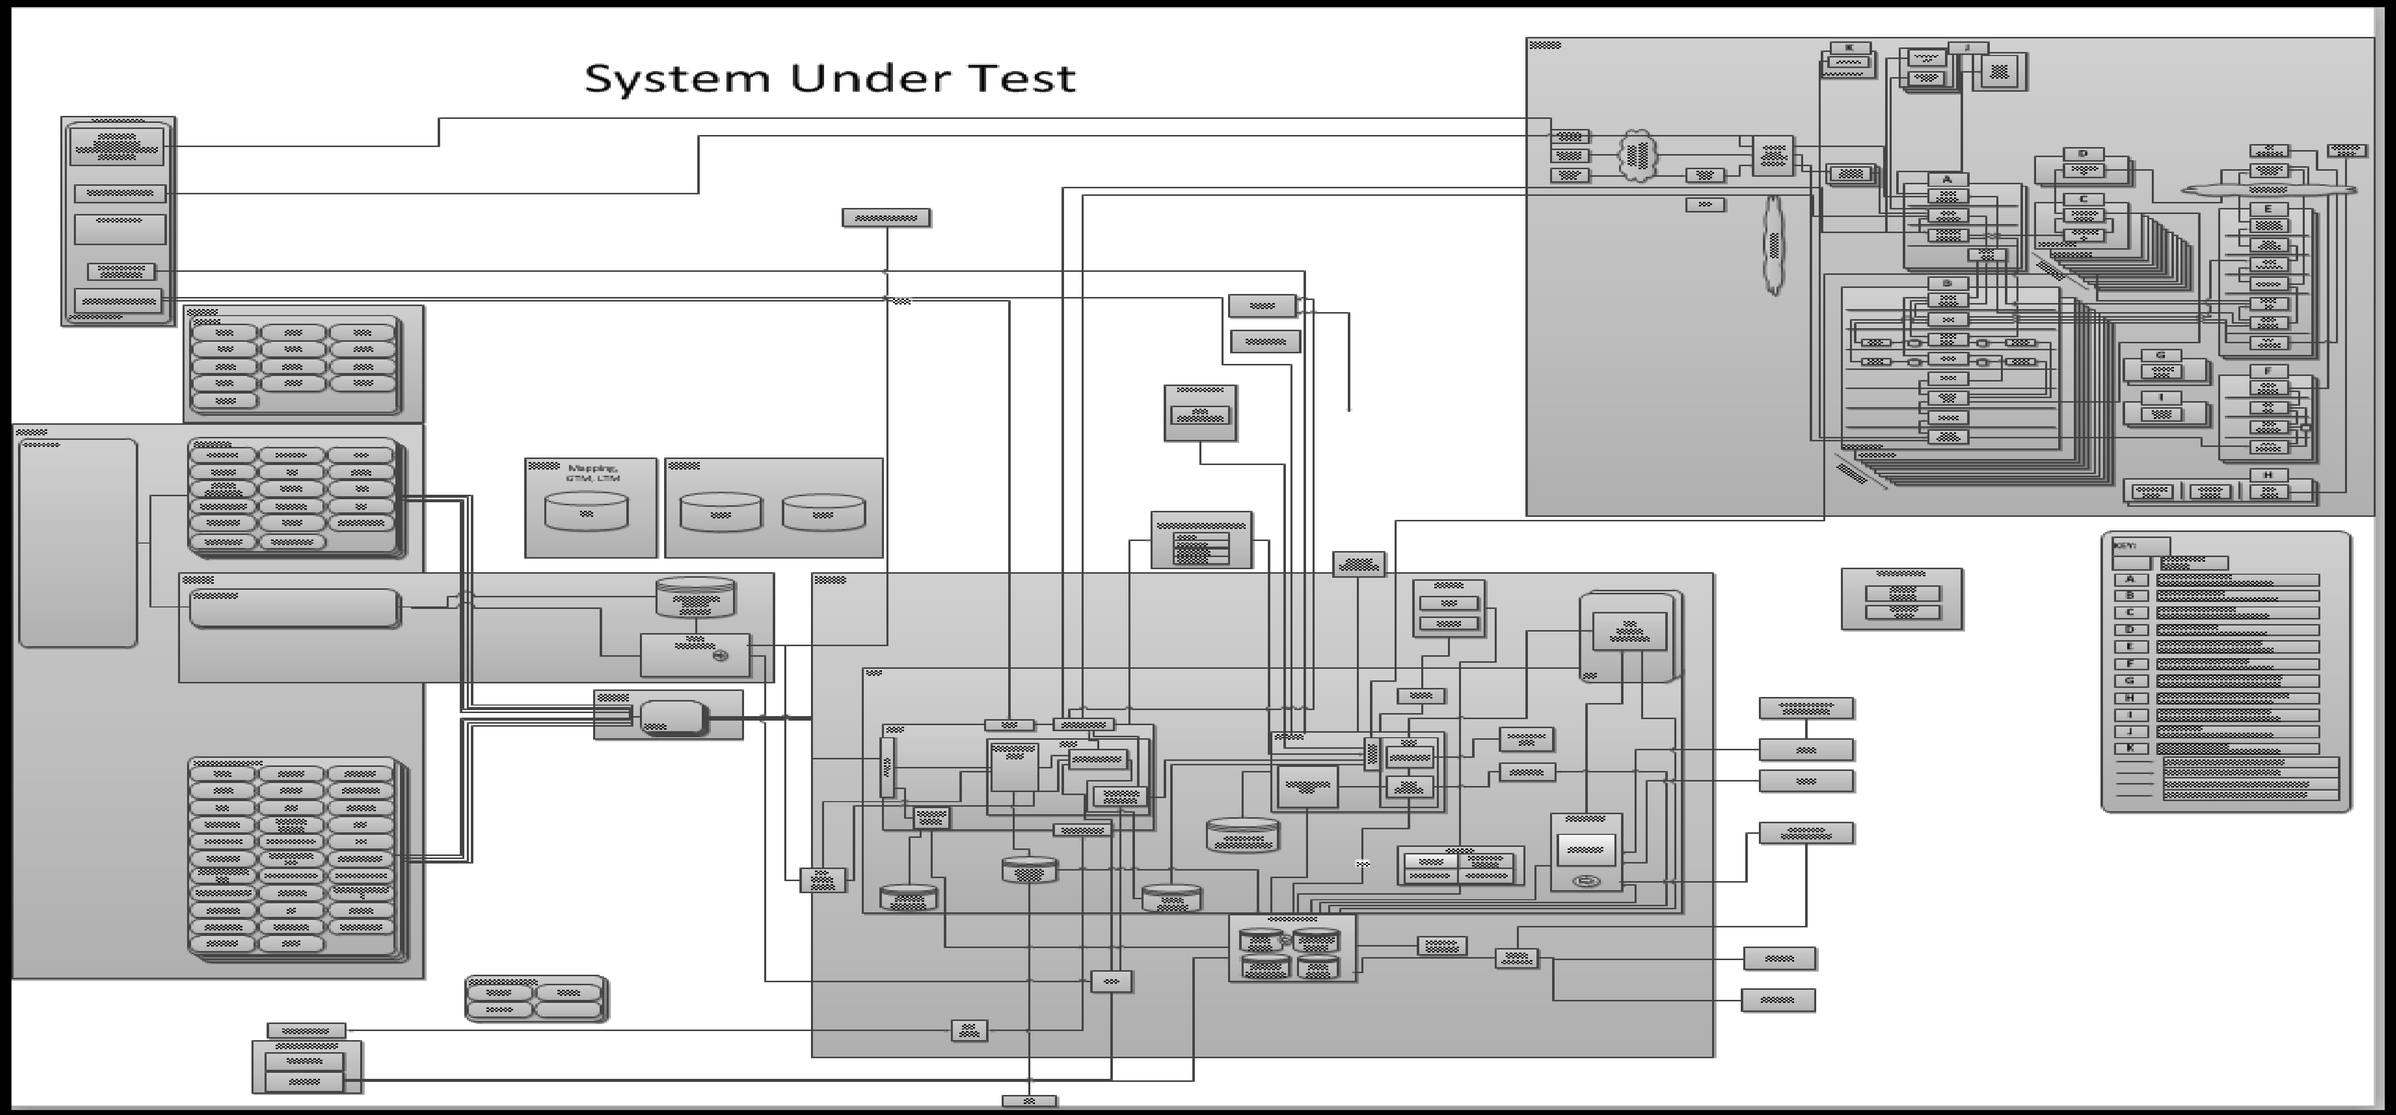

Ultimately, we settled on DR+prod

- DR provides a production scale mainframe with all data and applications
- Production is a “routed-out” segment of true production front-ends for a site
- Customers can be simulated from the cloud, getting closer to the real user experience
- We can performance test End-to-End!
- Tests simulate real production customer patterns and volumes on production systems
- Tests are repeatable!

Using the combination of DR+Prod creates opportunity

- Test at x times prod scale
- Replay production traffic
- Recreate windows of production activity
- Build repeatable test patterns/workload profiles
- Ensure veracity of tests through use of production data and application
- Break things with impunity – no messy cleanup!

The Process

So how do you make this happen?

Ultimately, it always boils down to the same three things…

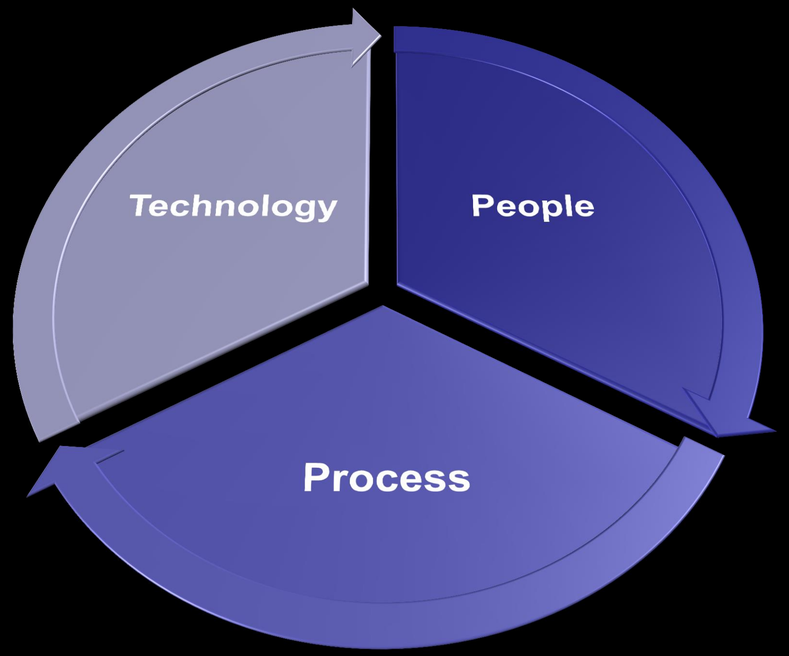

Unsolicited Observation – people are the hardest part…

First, we needed a team…

- Small number of specific resources
- Key characteristics
  - Firm knowledge
  - Process knowledge
  - Data knowledge
  - Executive presence
  - Assertiveness
  - Team oriented, but with independent spirit
- Fearless of failure

And an approach…no templates existed

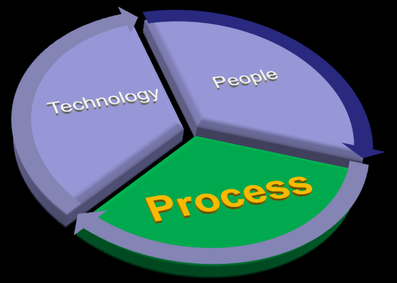

The complexity of the End-to-End testing requires that the team adopt a clear process for analyzing, planning, designing, and delivering the test models, data, and execution plans

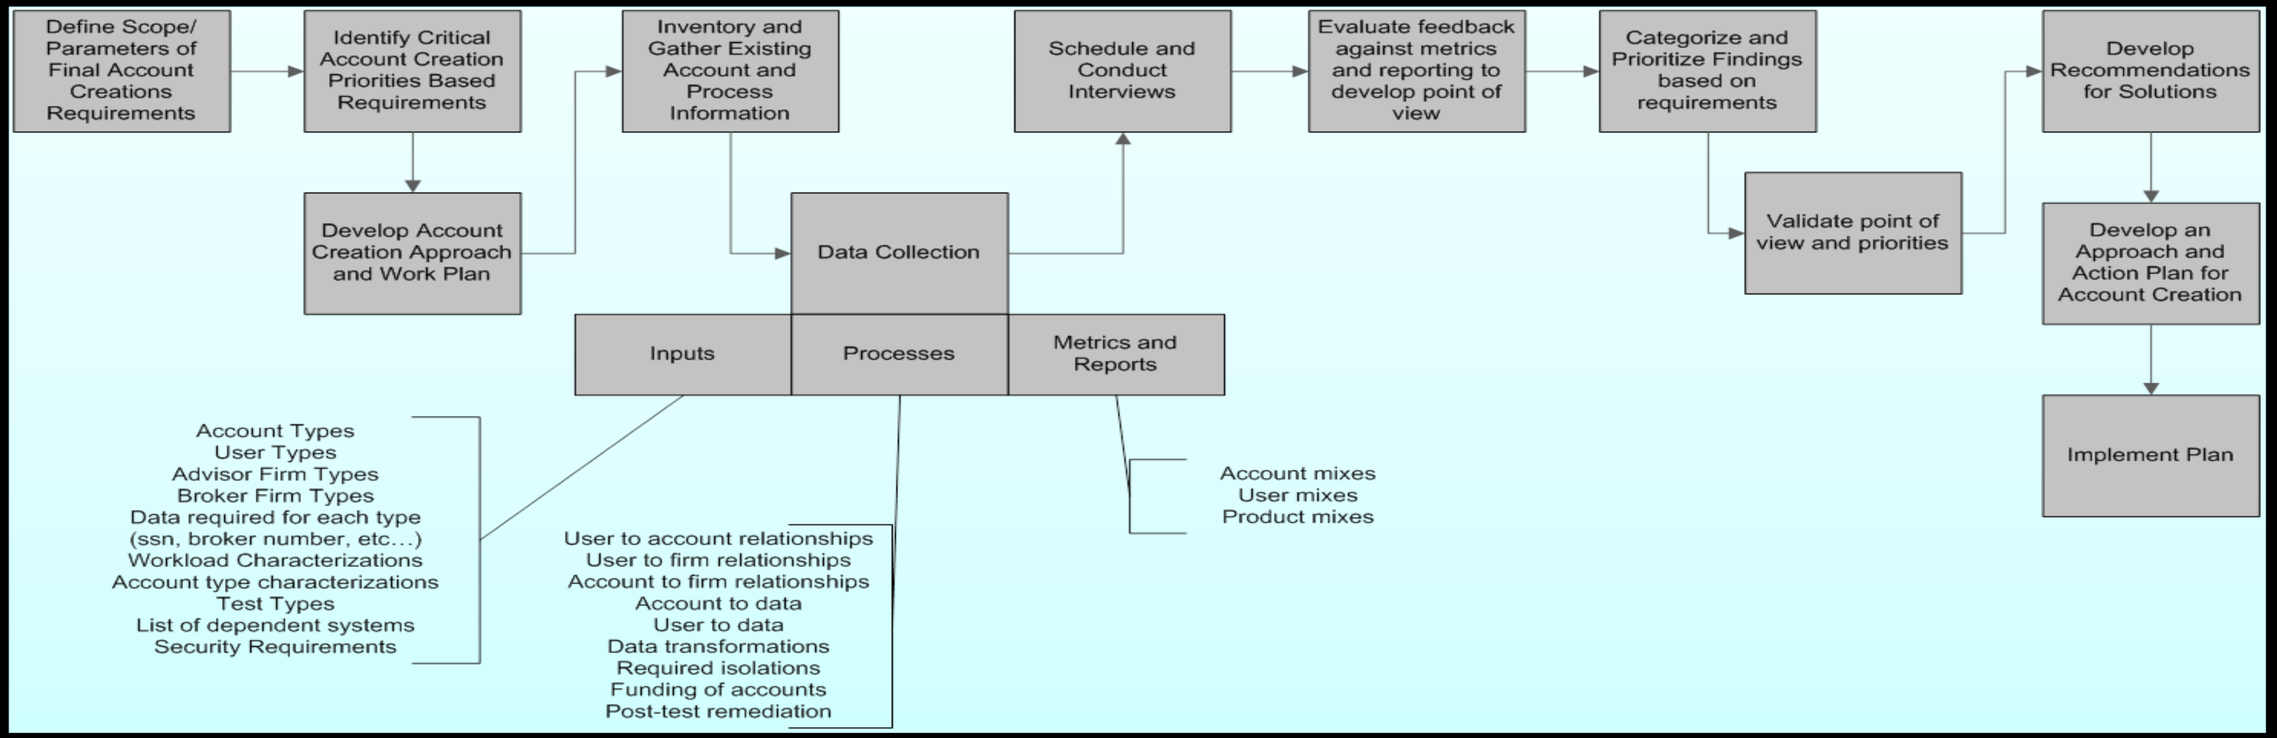

We needed to understand the systems and processes involved

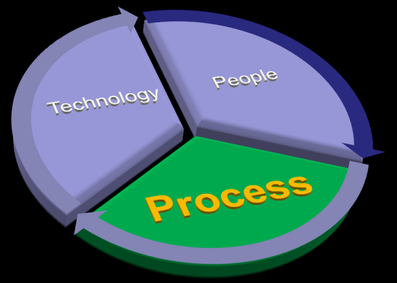

Brokerage systems are much more complex than the average application stack, and often have code and processes which have existed for decades.

Our testing requires deep process understanding

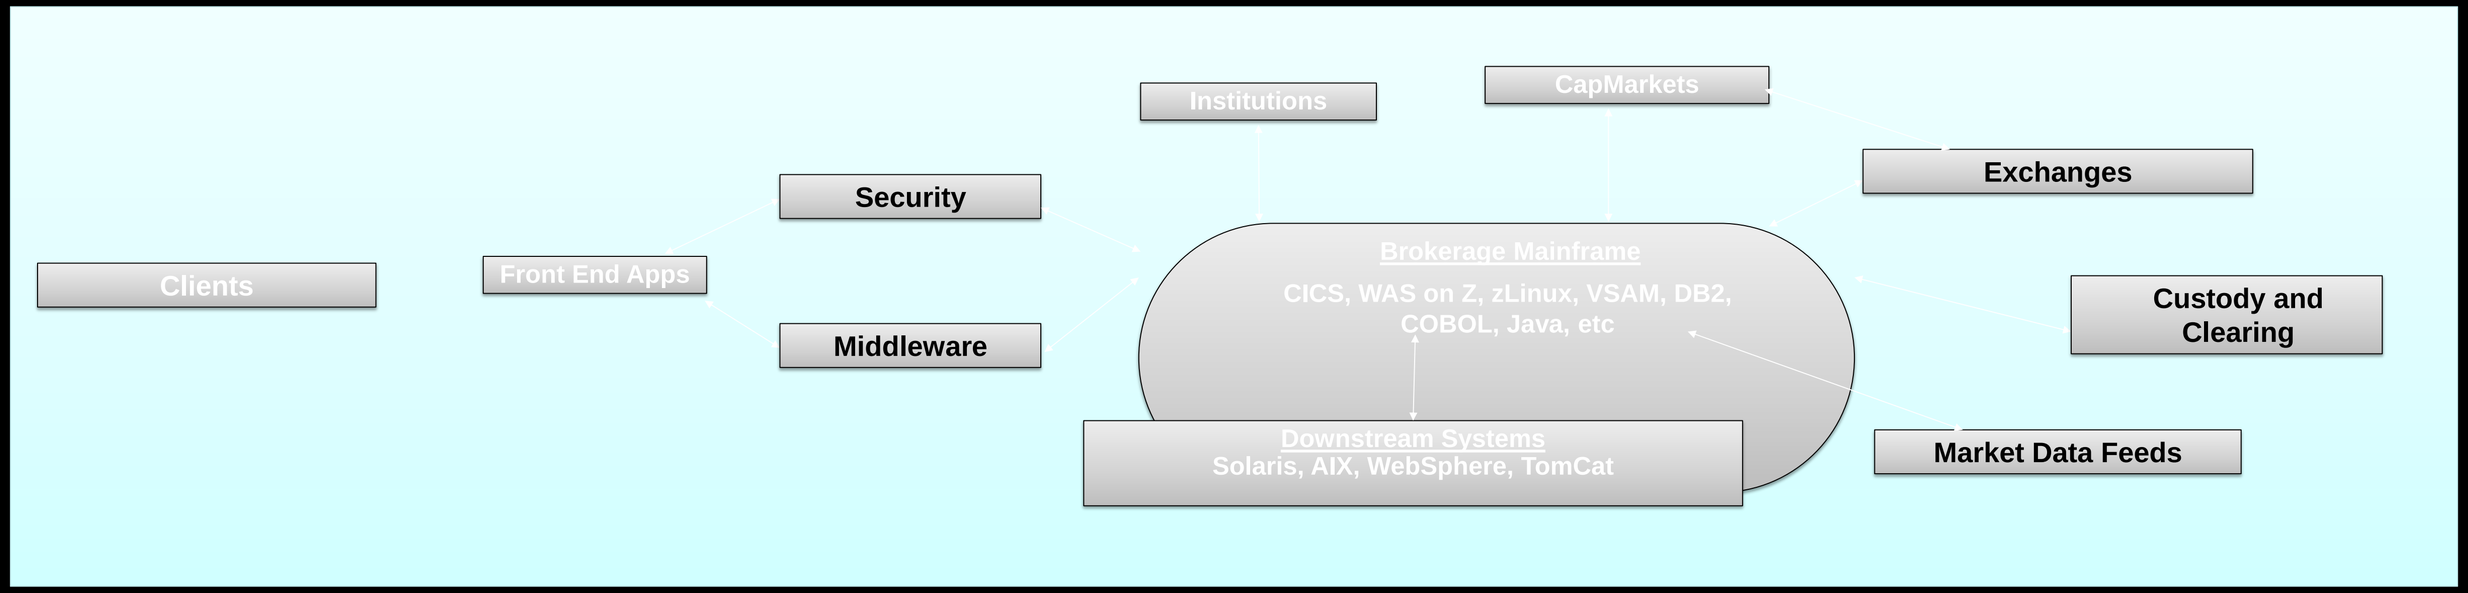

We needed to understand the data and entities

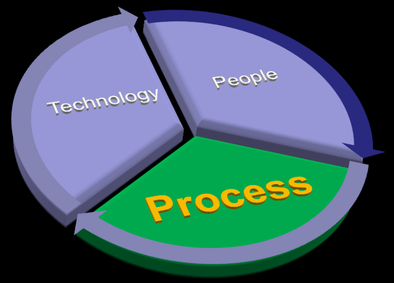

Security

- **Web Challenge Response**
  - PIN/UserID
  - Security questions
  - Password
  - Account mappings
- **3270 –**
  - UserID
  - password
  - TSO/Top Secret permissions

Internal Trade Routing

- **Product/Securities Info**
  - CUSIP
  - Symbol
  - Pricing
  - Update process
- **Corporate Actions**
  - CUSIP and Action
- **Streaming Data/Quotes**
  - Symbols/CUSIPS
    - Reuters
    - Activ Financial
    - Exchanges

Mainframe Applications

- **Order Management**
  - Account
  - Requestor id
  - Security
  - View
- **Customer Name and Address**
  - Account
  - Customer
- **Functional Applications**
  - View
- **Conditional orders**
  - Market data

Exchanges

- Symbols
- Prices
- Updates/changes

Institutions

Exchanges

Security

Clients

Front End Apps

Middleware

Brokerage Mainframe
CICS, WAS on Z, zLinux, VSAM, DB2, COBOL, Java, etc

Custody and Clearing

User Types

- **Retail**
  - .Com - user ID pass
  - Frequest Traders – userID pass
  - Others - ??
- **Brokers/Dealers**
  - Broker – userID pass
    - Administrator
  - Client – userID pass
- **Advisors**
  - Channel1 – userID pass
  - Channel2 – userID pass
- **3270 Manual** – userID pass

Middleware

- Account number
- Requestor ID
- View Request
- Physical Routing

Downstream Systems
Solaris, AIX, WebSphere, TomCat

Market Data Feeds

Market Data

- Streaming
  - Quotes
  - News/Charts
- Alerts
  - Updates to users

Clearing

- Account number
- Intraday positions

While always assessing risk

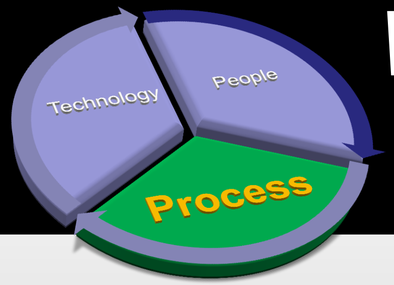

When testing in Production…

Risk management is KEY…

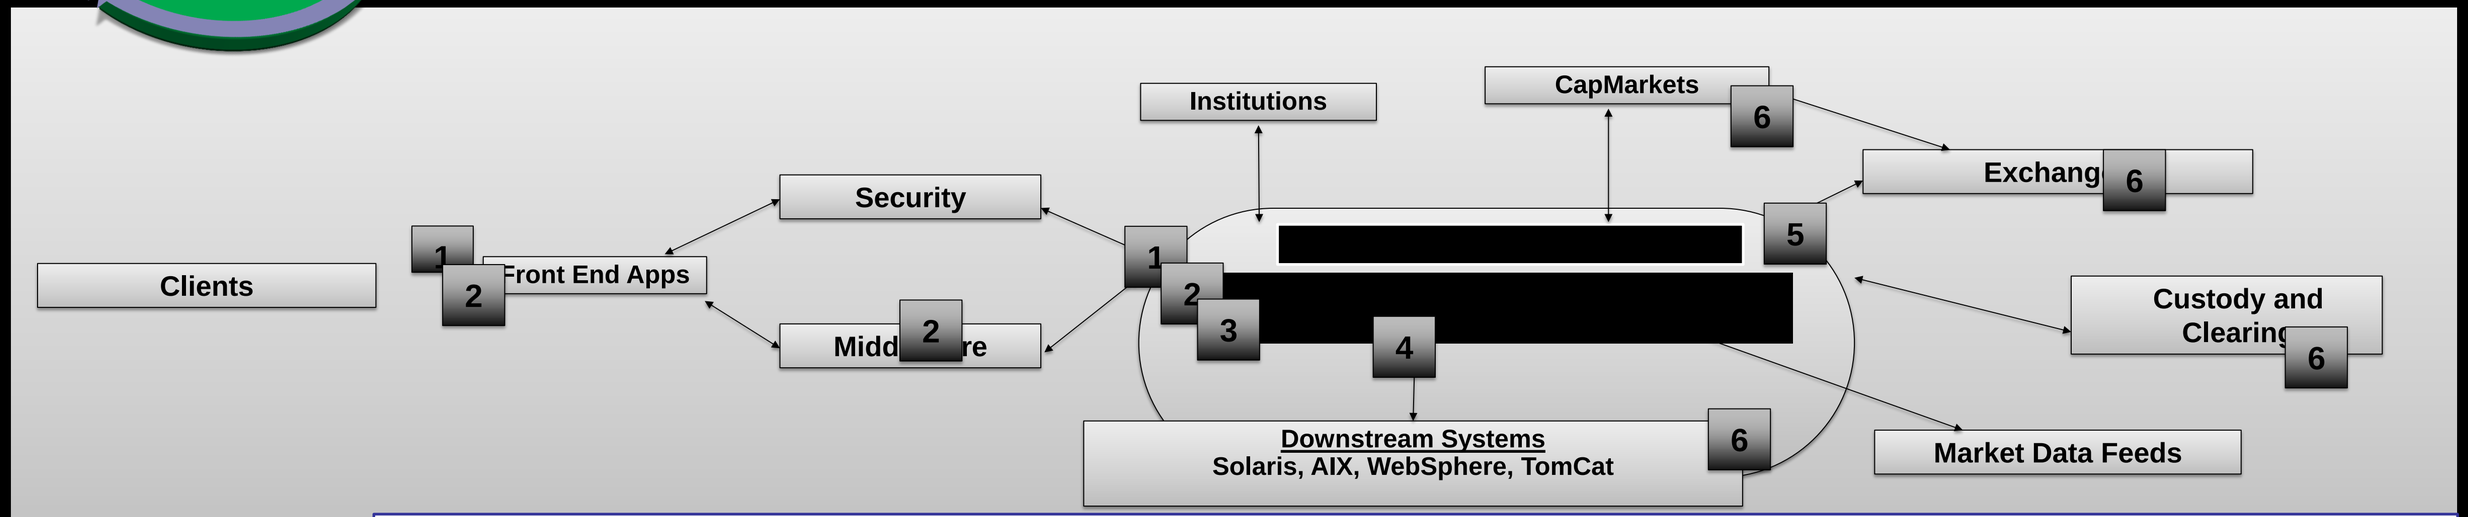

1. Preparation of transaction loads for the load generation scripts will require that all transactions are submitted to clearly identified securities which are designated as test only securities
2. Physical and system access to the load generation data and facilities will be restricted to team involved directly in the test
3. Test transactions will only be submitted through users and into accounts which are clearly identified as test only users and accounts
4. Test users and test accounts will only be attached to specific branch codes designated as test only branches
5. The only securities which will be included in any transaction during test are those securities created/assigned specifically to the test
6. External reporting and financial impacts will be identified and isolated, and tested prior to test

Navigating the enterprise presents challenges

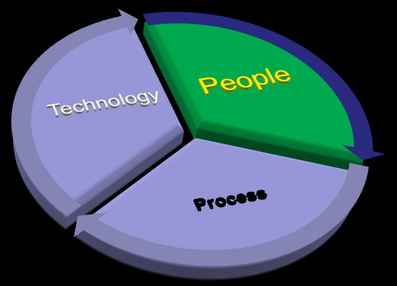

- Three different brokerage groups and all of their friends
  - Technology
  - Account Setup
  - Firm Setup Team
  - Deployment
  - Product Management
  - Risk and compliance
  - Functional, Integration, and Performance QA
- Corporate Risk and Compliance
- Clearing Services
- Order Management
- Name and Address

- Capital Markets
- Security Master/Product Master
- Fees and Commissions
- Electronic Notifications
- Corporate Actions
- Streaming Data and News
- Market Data
- Operations
- Printing and reporting
- Offshore QA Teams
- Purchase and sale/Books and Records
- Security
- Third Party providers

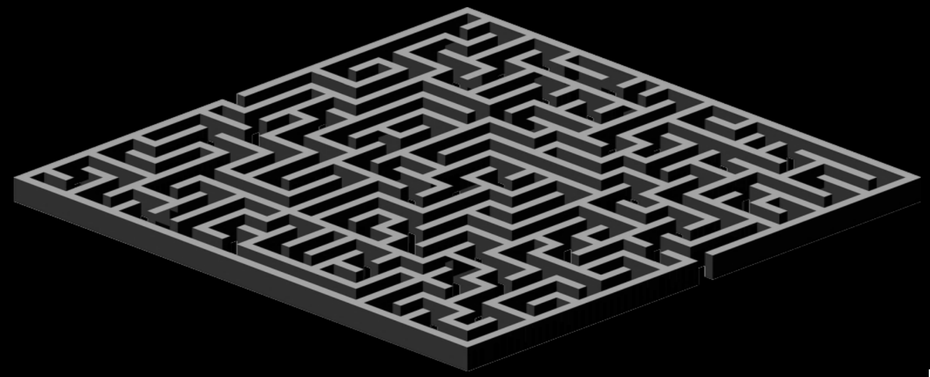

We needed help from a lot of people…

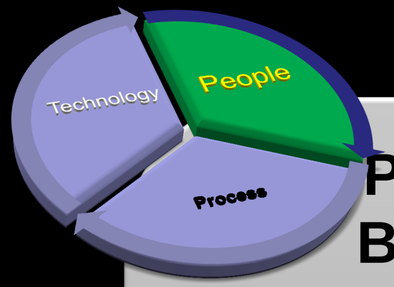

Participation by Business Group

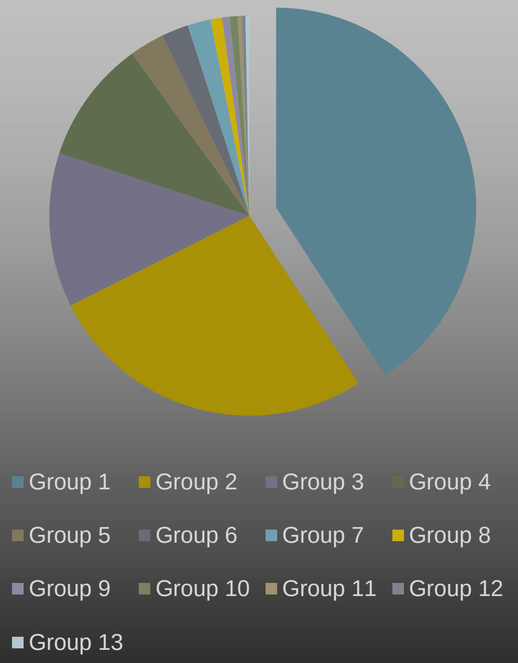

Participants by Division

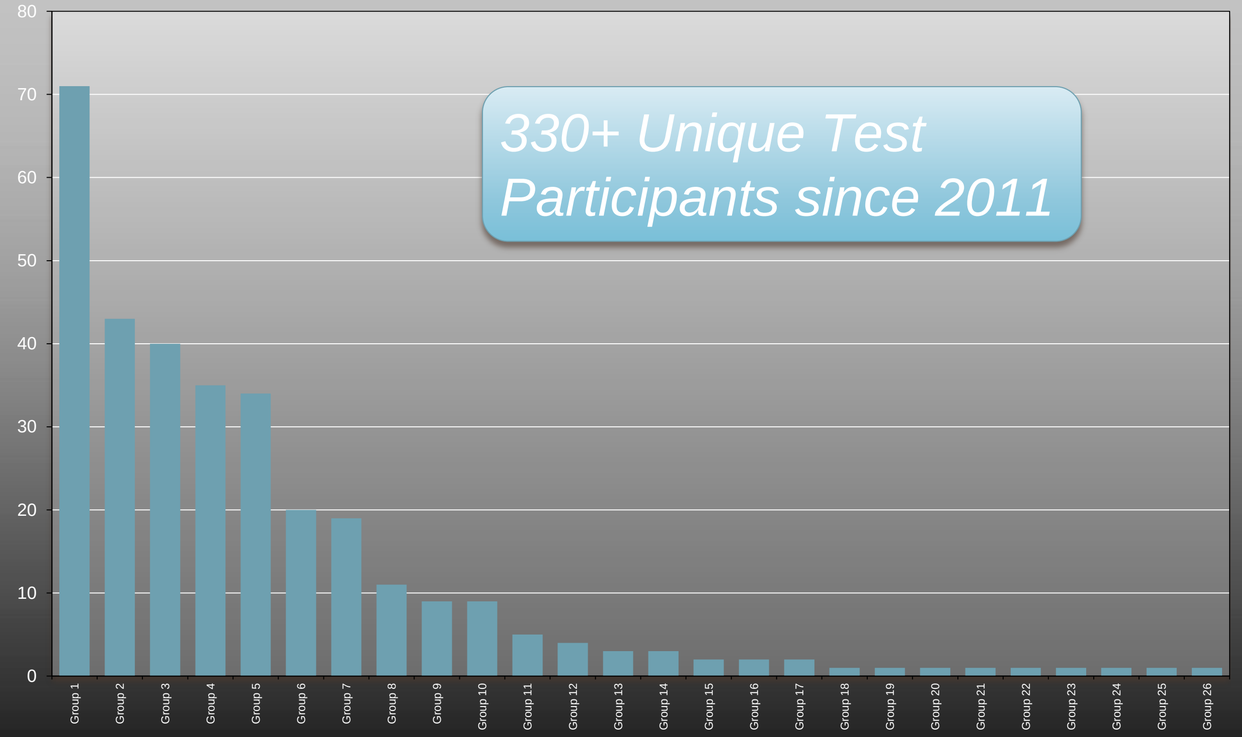

CMG Computer Measurement Group

Building the test

Systems are the vehicle, but data is the fuel…

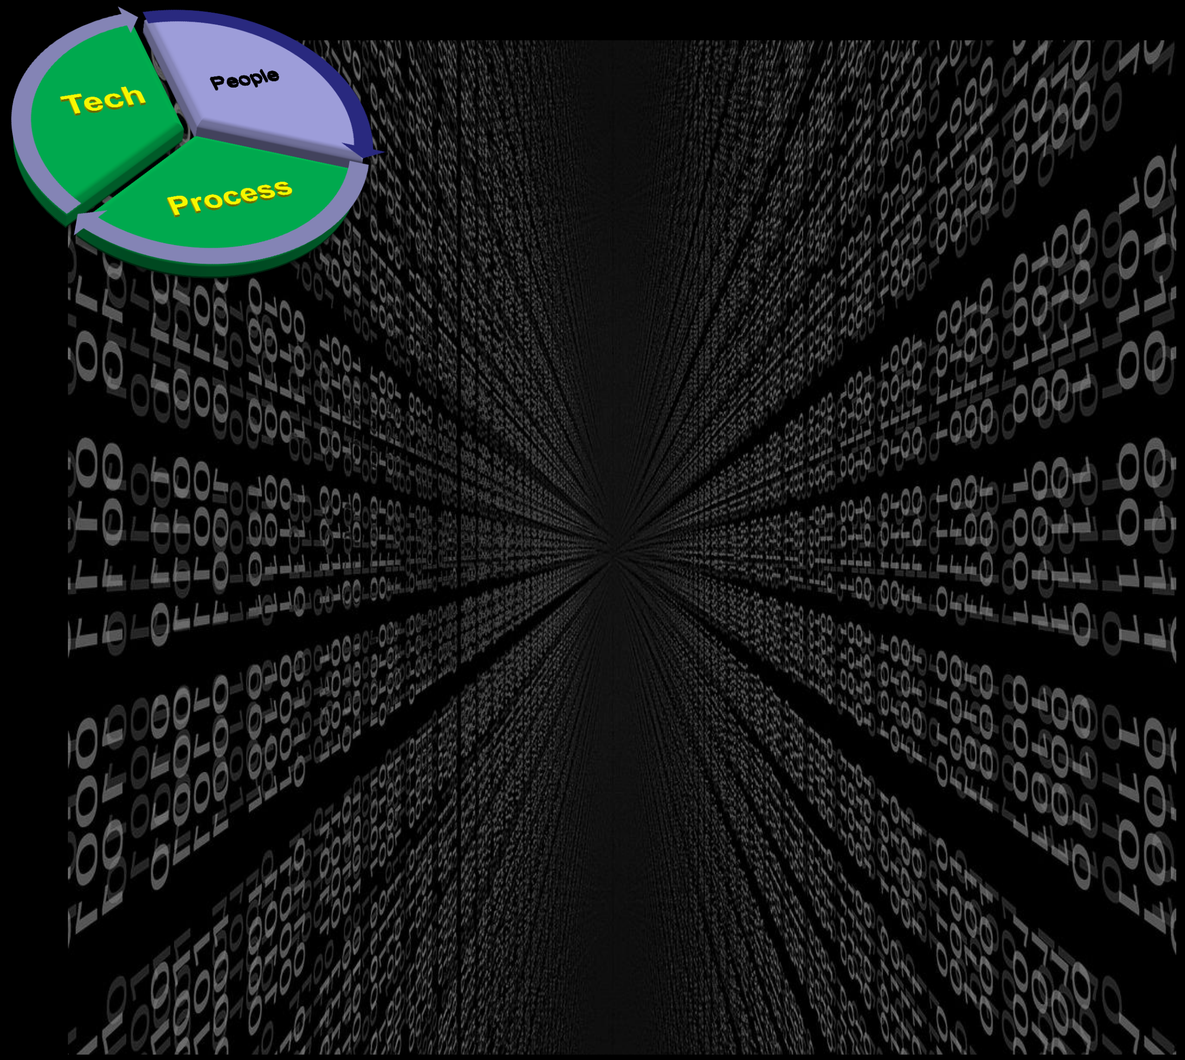

- Approximately 110,000 user IDs
- Approximately 300,000 investment accounts
- Synthetic marketplaces
- UserID and account combinations to cover multiple business channels
- Full coverage of all registration/account types
- Creation of very diverse account holdings

Data needs to be accurate and SAFE, so masking DR is critical

DR is a copy of real customer data

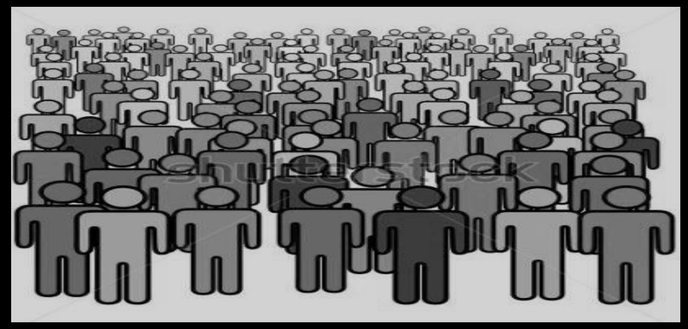

We use that, +masking to deliver realistic, YET Safe, users and accounts

PRODUCTION	Portfolio $	Acct	Pos	Trades/yr	Name	Address
RealCustID1	$1,097,230	5	36	220	John Doe	30 Elm St
Real						
Real						
Real						
x10						

DR	Portfolio $	Acct	Pos	Trades/yr	Name	Address
RealCustID1	$1,097,230	5	36	220	John Doe	30 Elm St
RealCustID2	$10,234,400	22	135	600	Jane Doe	20 Maple St
RealCustID3	$65,340	2	12	10	Tom Smith	40 Main St
RealCustID4	$113,434	2	5	32	Sue Jones	10 River St
x100,000	etc.	…	…	…	…	…

Masking batch jobs

DR	Portfolio $	Acct	Pos	Trades/yr	Name	Address
FakeUser1	$10	1	1	0	S3PEED 1	1 SPEED ST
FakeUser2	$10	1	1	0	S3PEED 2	1 SPEED ST
FakeUser3	$10	1	1	0	S3PEED 3	1 SPEED ST
FakeUser4	$10	1	1	0	S3PEED 4	1 SPEED ST
x100,000	$10	1	1	0	…	…

CMG Computer Measurement Group

“Splicing” allows us to realistically look like our customers

Real DR copies real customer data, so we wipe it

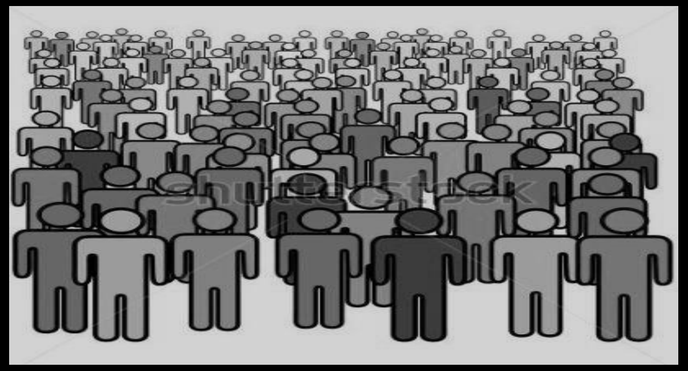

But even fake users need liquidity, so we provide it

Masking batch jobs

DR	Portfolio $	Acct	Pos	Trades/yr	Name	Address
FakeUser1	$10	1	1	0	S3PEED 1	1 SPEED ST
FakeUser2	$10	1	1	0	S3PEED 2	1 SPEED ST
FakeUser3	$10	1	1	0	S3PEED 3	1 SPEED ST
FakeUser4	$10	1	1	0	S3PEED 4	1 SPEED ST
x100,000	$10	1	1	0	…	…

“splicing” batch job assigns cash and positions to fake users

DR	Portfolio $	Acct	Pos	Trades/yr	Name	Address
FakeUser1	$1,097,230	5	36	220	S3PEED 1	1 SPEED ST
FakeUser2	$10,234,400	22	135	600	S3PEED 2	1 SPEED ST
FakeUser3	$65,340	2	12	10	S3PEED 3	1 SPEED ST
FakeUser4	$113,434	2	5	32	S3PEED 4	1 SPEED ST
x100,000	etc.	…	…	…	…	…

Additional Key Processes

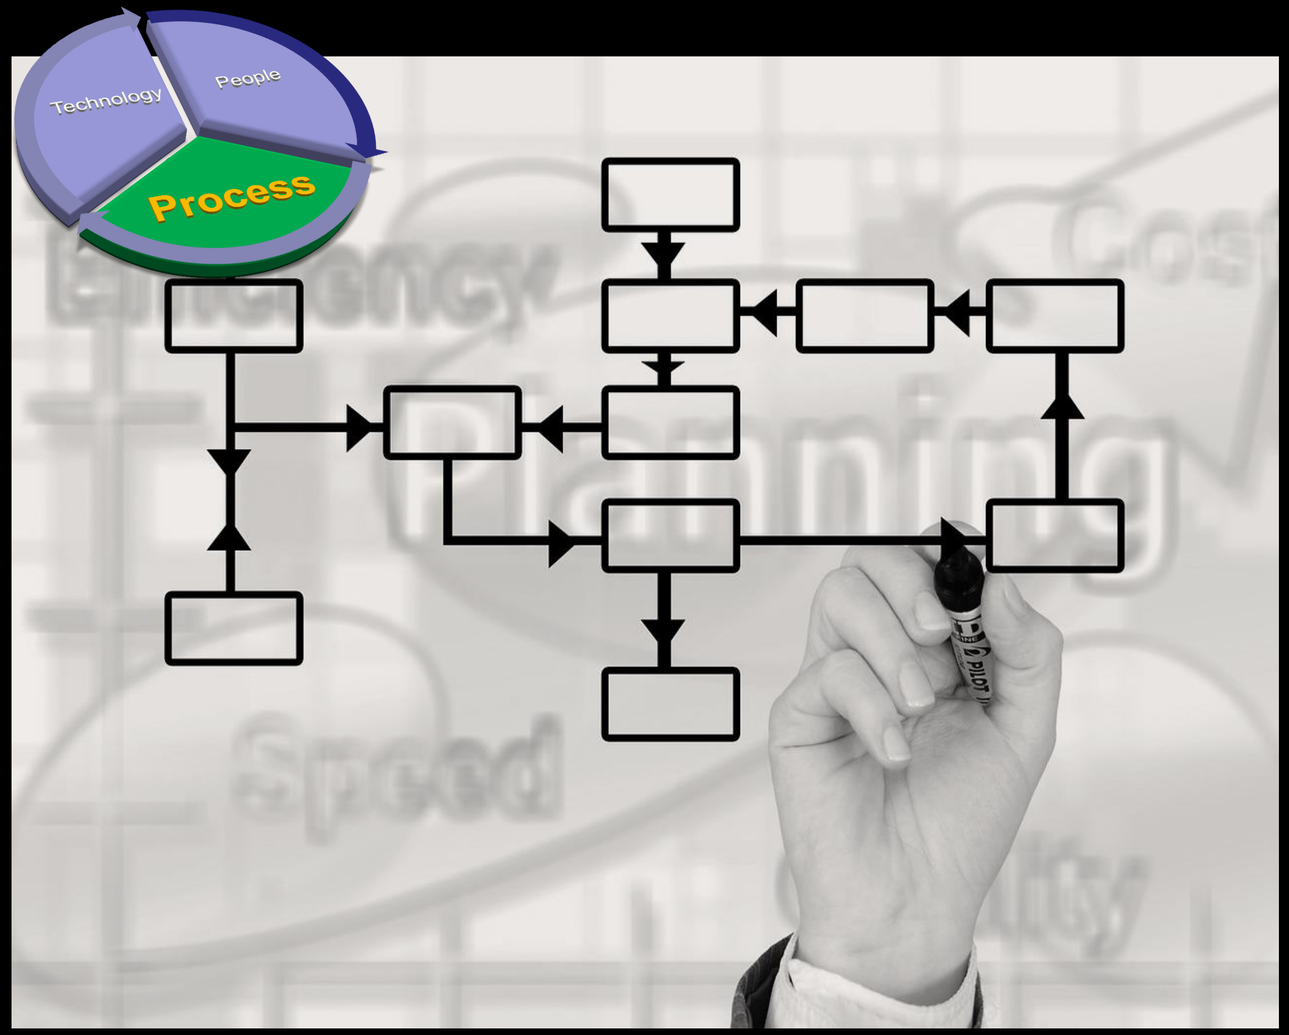

- Capture replay data
- Process data to the load drivers
- Establish all security perimeters
- Build the DR environment

Running the test

The final test platform integrates Prod, DR, and Test…

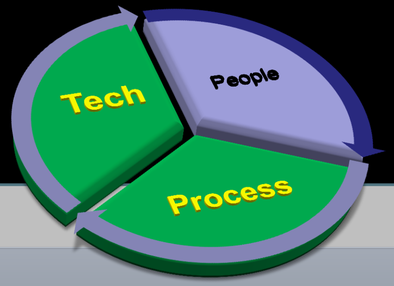

Institutions

CapMarkets

Exchanges

Security

Clients

Front End Apps

Middleware

Brokerage Mainframe
CICS, WAS on Z, zLinux, VSAM, DB2, COBOL, Java, etc

Custody and Clearing

Downstream Systems
Solaris, AIX, WebSphere, TomCat

Market Data Feeds

We can test in “production” and never impact customers

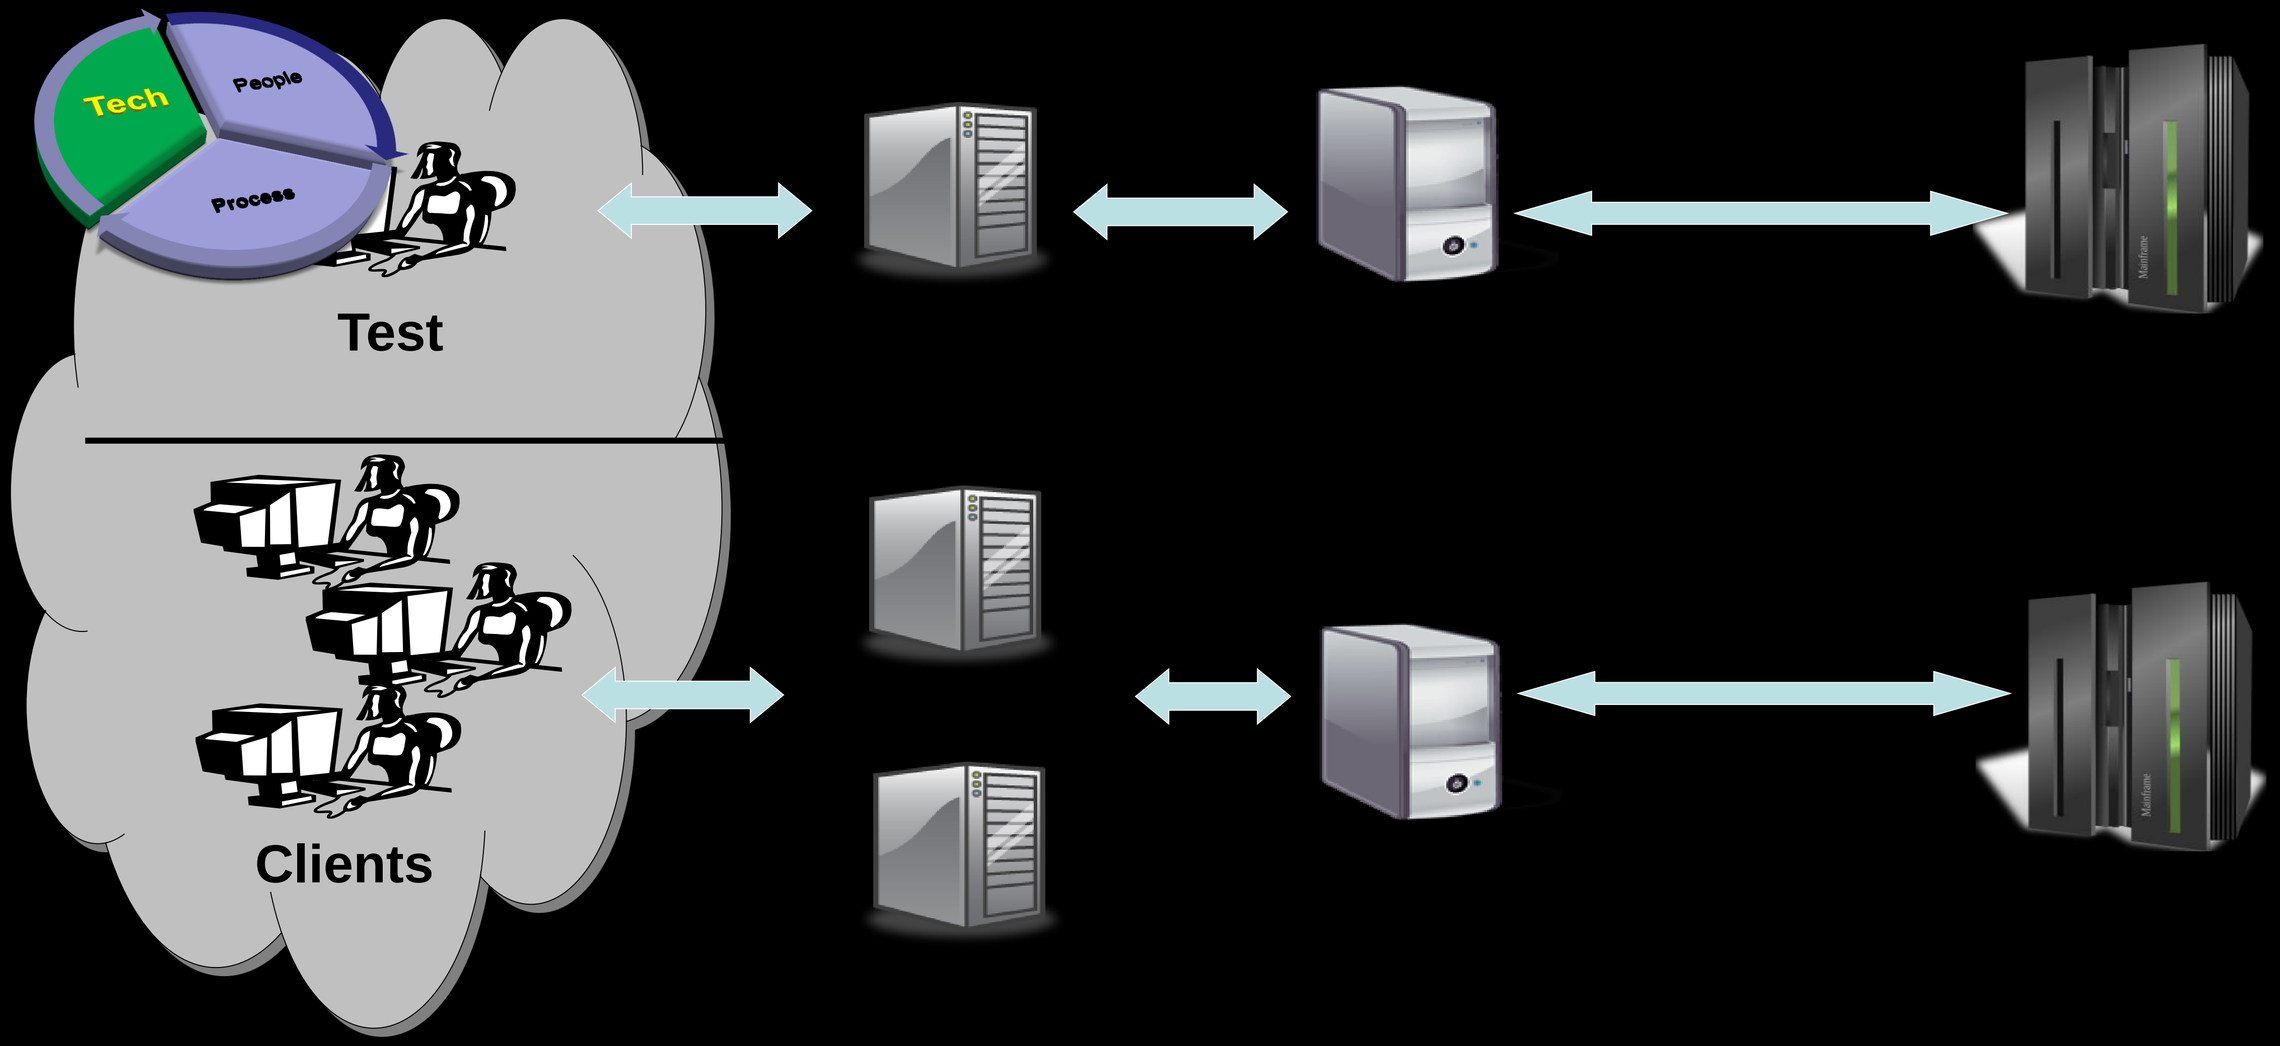

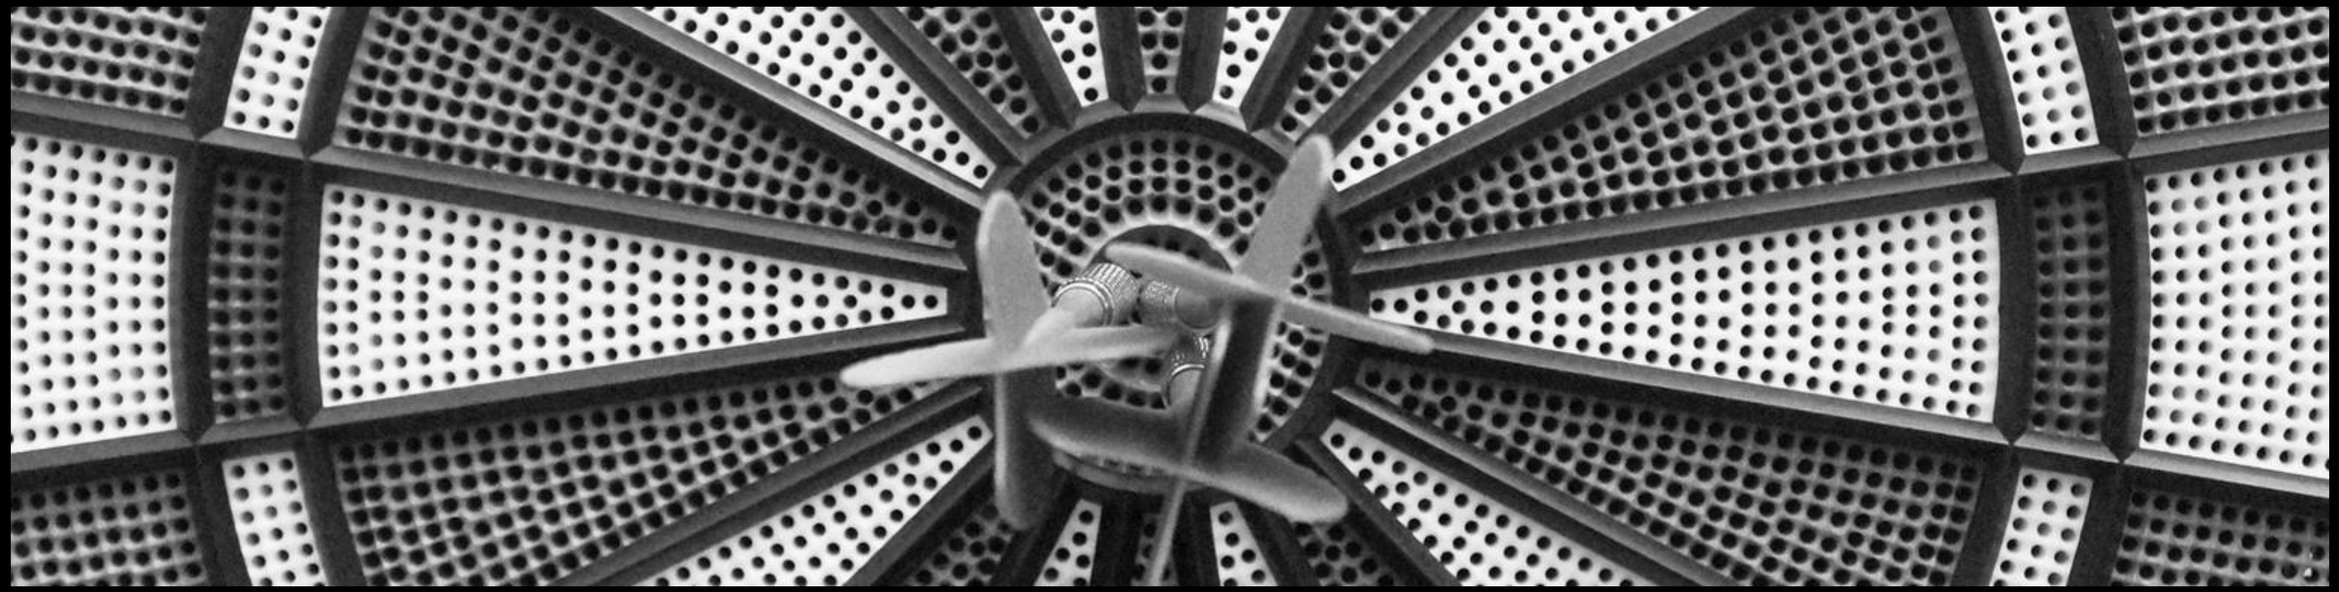

The Results

And testing leads to fixes before our customers find them…

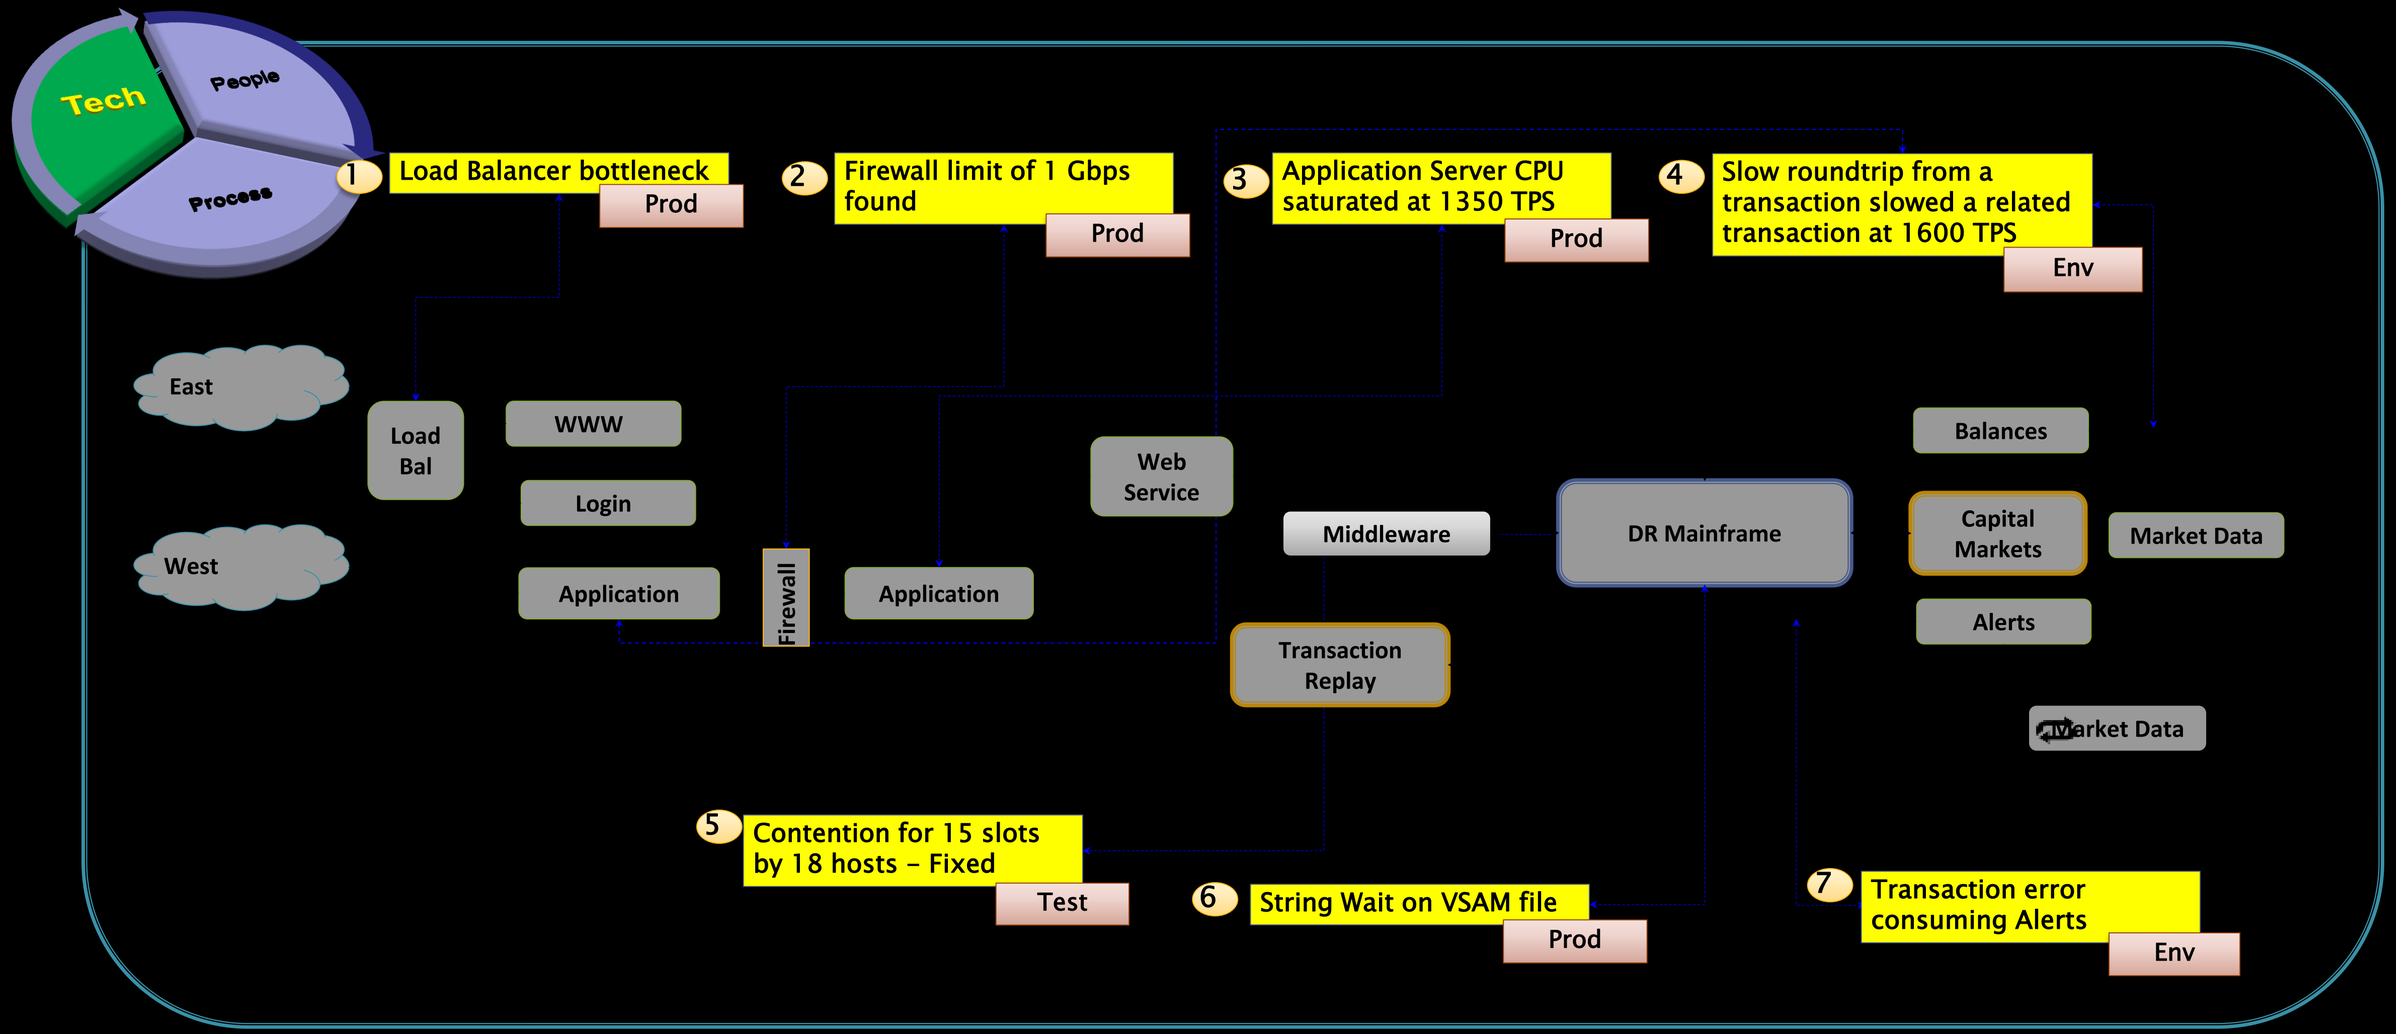

And because of these tests, there has been much joy!!

The Ultimate Result:
We quickly adapt to changing conditions with relatives ease and know our systems limits.

But there is always more to do, as management is still asking…

- When are you going to…
  - Build pre-prod testing
  - Design selectable workloads
  - Improve infrastructure alignment
  - Include additional channels
  - And a bunch of things that they have not yet imagined

Takeaways

> “*It is not the strongest of the species that survives, nor the most intelligent. It is the one that is most adaptable to change*“. *Leo Megginson*

- Testing like this takes patience. It is a journey, and a long one at that.
  - This is not a “I wrote a script and ran a test” scenario
- Relationships are critical
  - To be successful, you need help from a lot of people
- You need to be ok with failure
  - In complex testing, with dynamic requirements and targets, you will fail
    - The upside – you will learn an amazing amount
- There are many ways to “skin a cat”
  - Be creative – Stretch your imagination – Try to have fun!

CMG Computer Measurement Group

Questions?

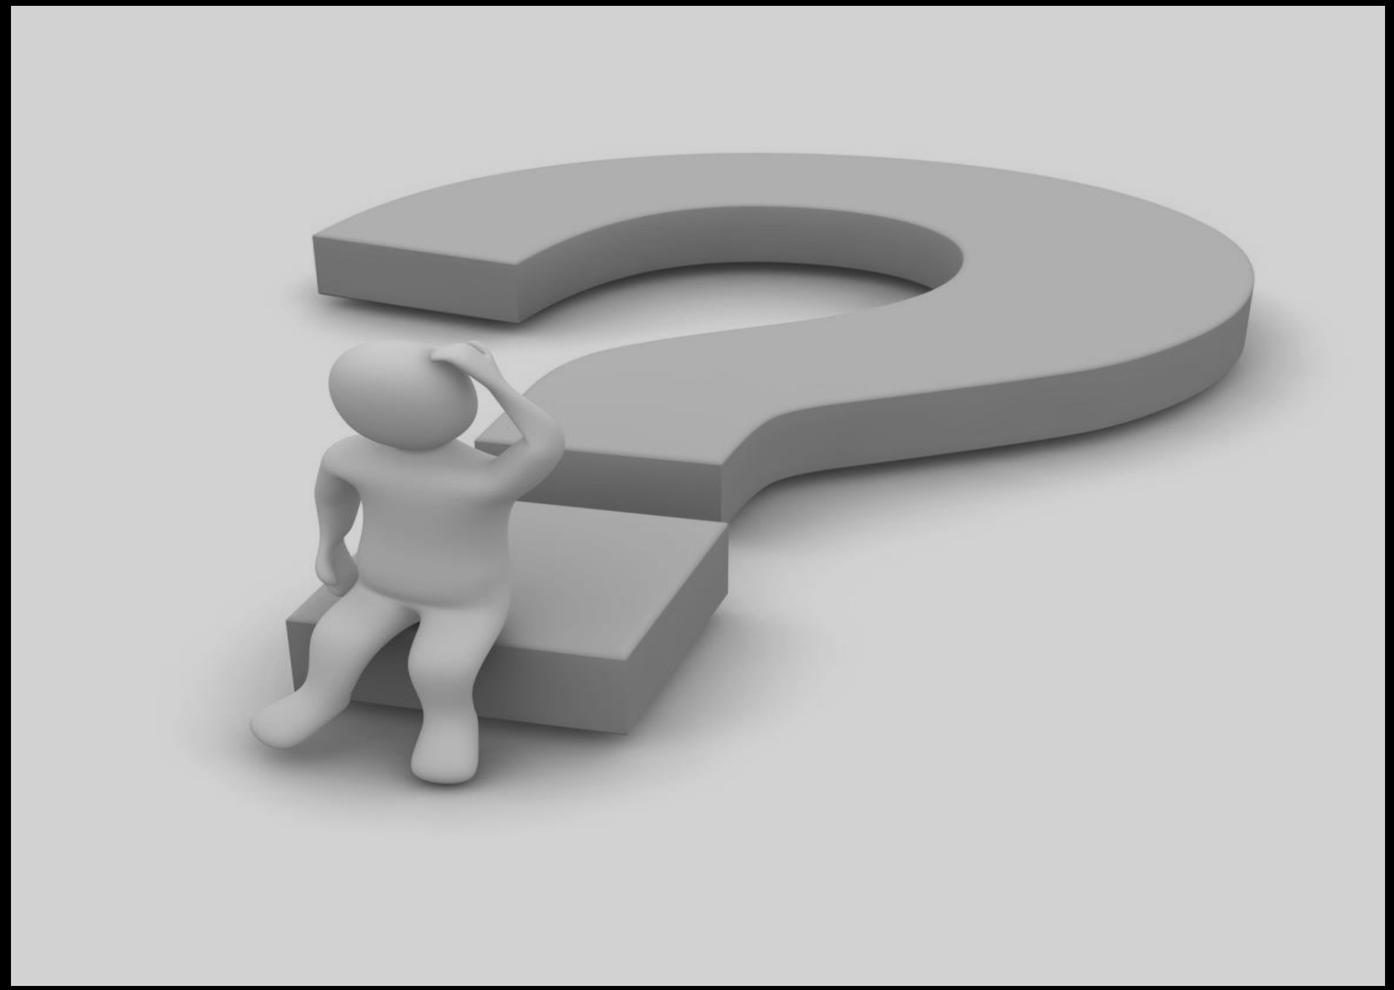

Thank you!