

John Davies
CTO

Incept⁵



In-Memory Message & Trade Repositories

QCon - London
8th March 2013

I really don't like

- Writing slides
- I try to talk independently from the slides
- Yesterday there was a great keynote talk on minimising slides
- Sadly I missed it because I was writing my slides :-(
- So today might just be my last slide deck!

I do like

- Talking at conferences
- Meeting all the interesting people
- Sharing my experiences
- And drinking beer

Agenda

- External influences on the banks
- The scale of the changes
- The scale of complexity
- Just chuck it into a database with ORM! - NOT!
- In-Memory
 - Perhaps a demo

Dodd-Frank Wall Street Reform

- Trying to keep this to 2 slides...
- Most of the financial world got a nasty shock in 2008
 - Things started going wrong in 2007 though
- There was no single cause but it's fair to say risk management got a little out of hand
- In July 2010 congress passed Chris & Barney's Dodd-Frank Wall Street reform
 - A week later President Obama signs it into law
- October 2011 the European Commission makes similar proposals



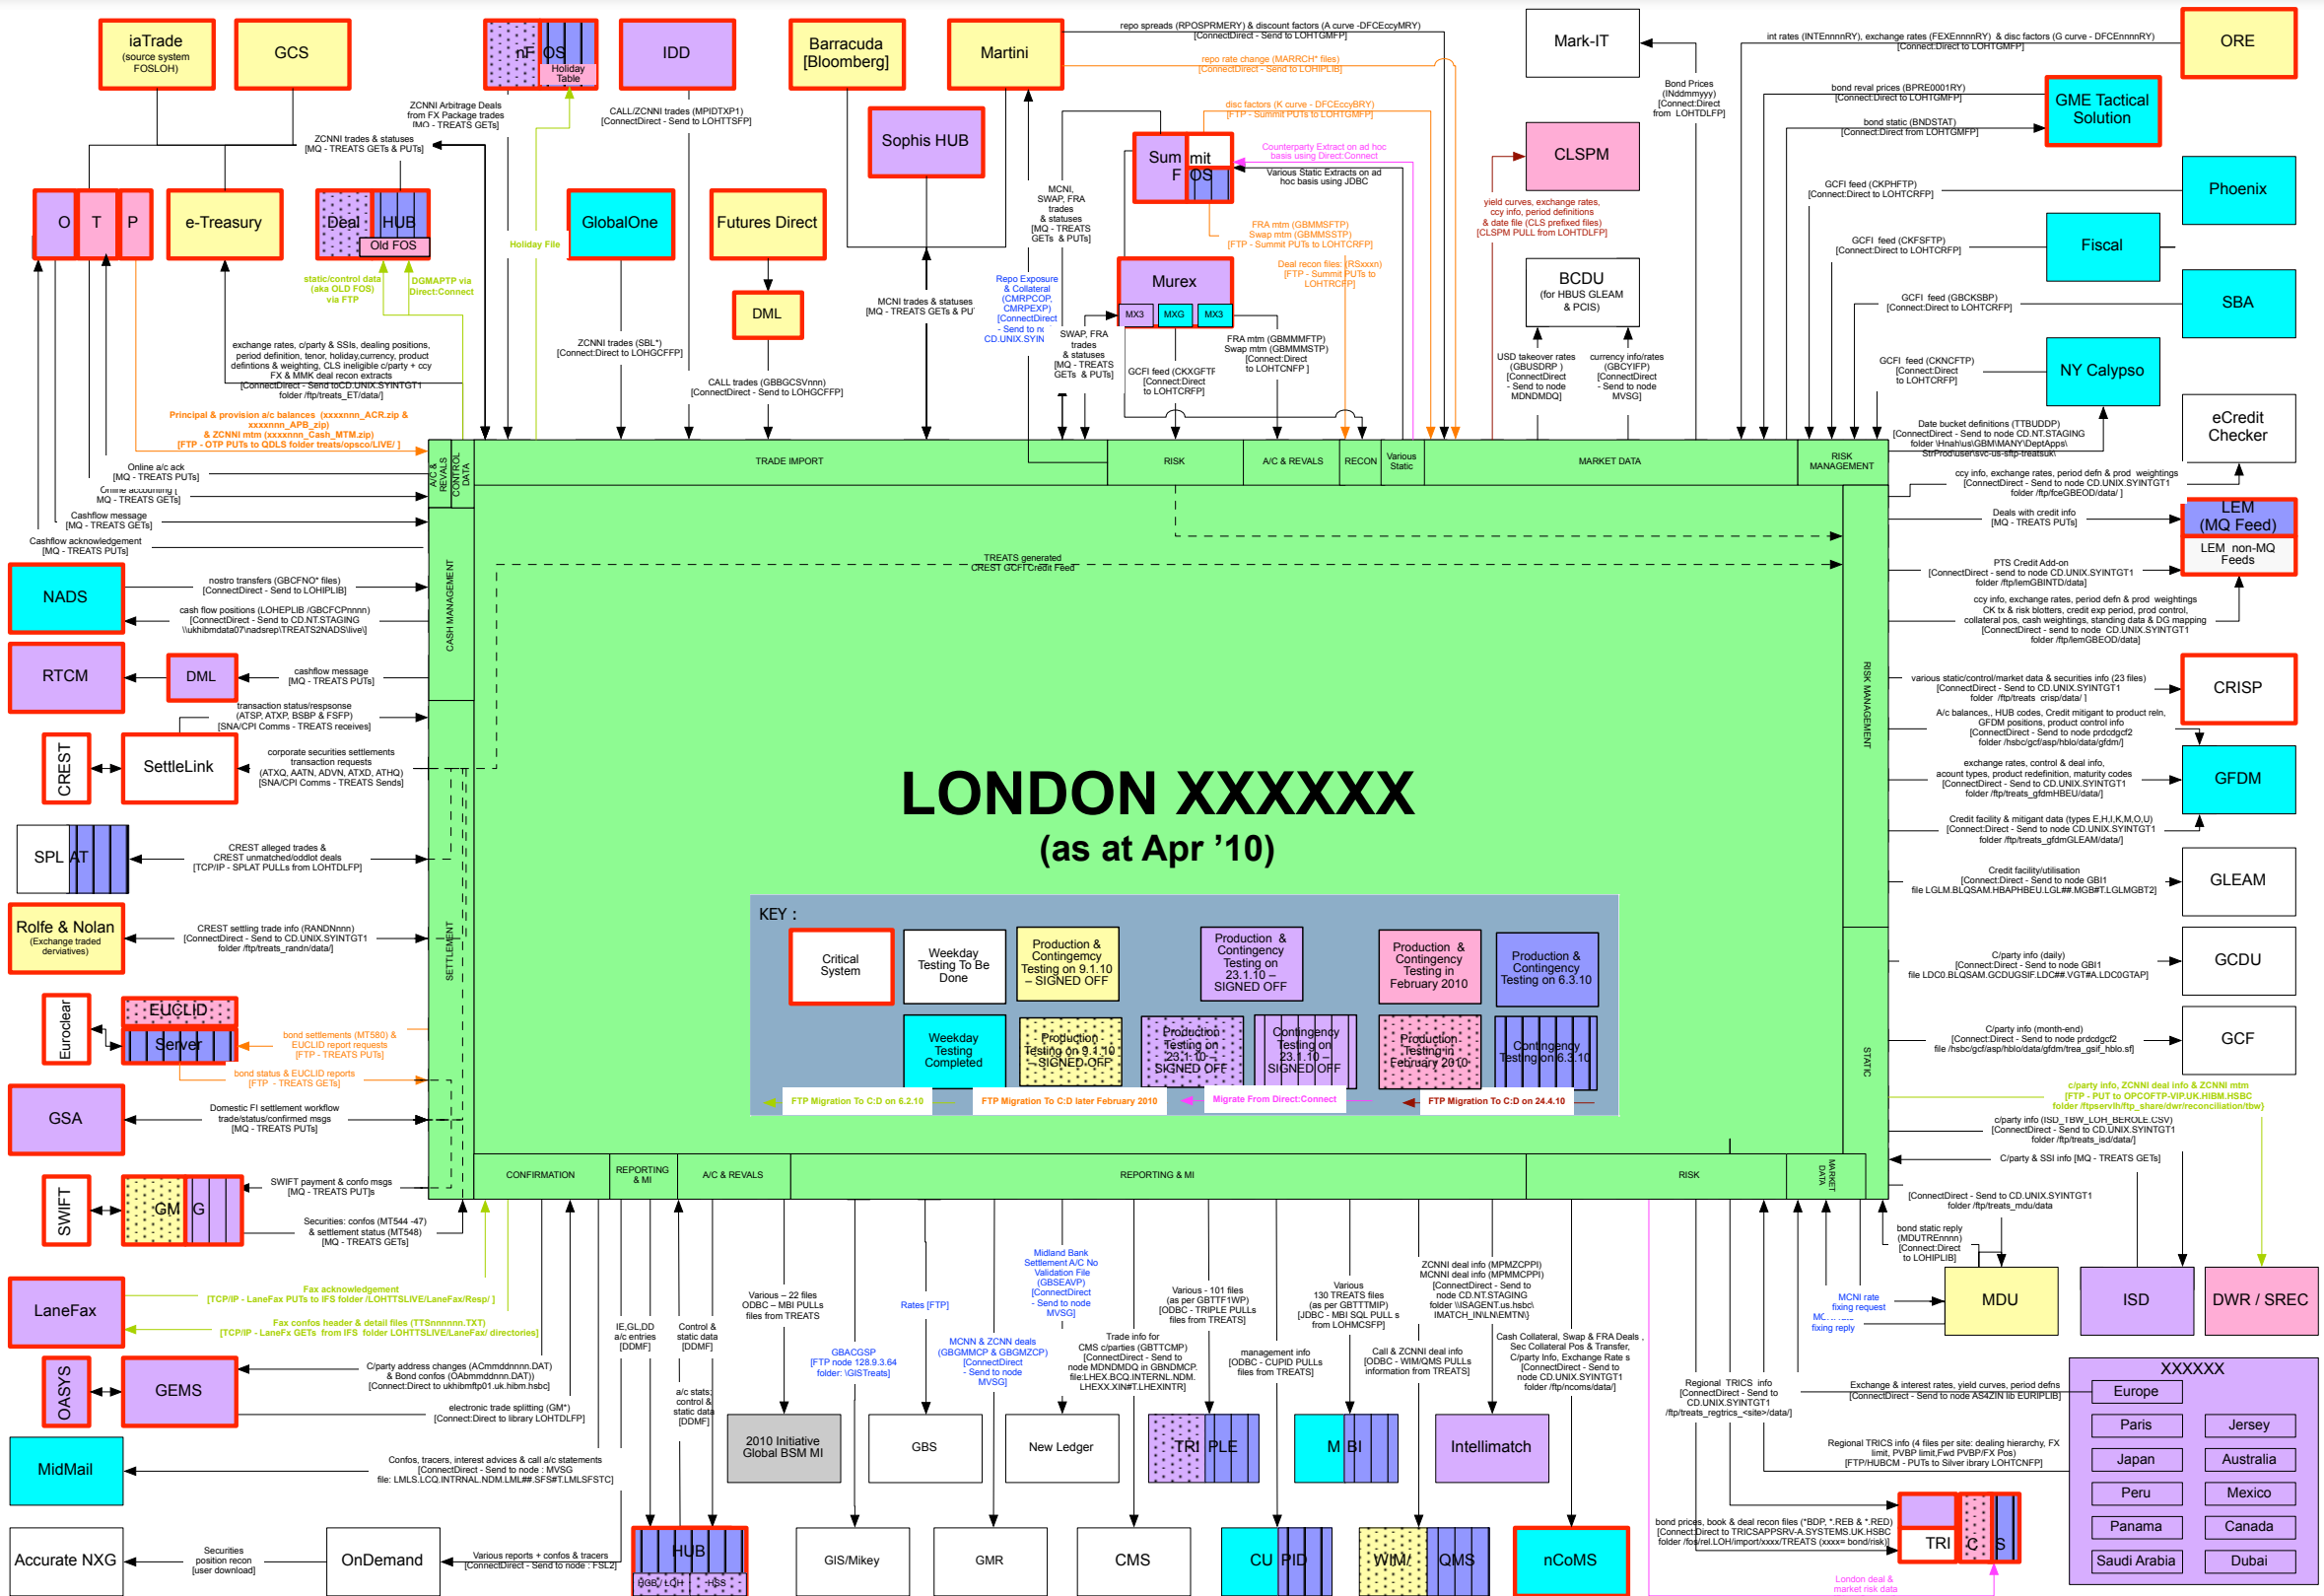
So what's happening?

- The Dodd-Frank act is wide ranging, the impact varies by geography and business domain
 - broker/dealers, asset managers, hedge funds in the OTC market
- It is already live in the US (12th Oct 2012), EMEA reporting begins in early/mid-2013
- The bodies (above) have to register certain types of Swap (the majority) with a central Swap Data Repository (SDR)
 - The DTCC is one of these SDRs
- Basically then all these financial institutions have a lot of work to do to comply

OK just one more slide...

- Dodd-Frank is not the only catalyst for change this year...
- SWIFTNet Derivatives (FpML over SWIFT)
- Accelerated ISO-20022 adoption
 - DTCC Corporate Actions, JASDEC and T2S
- And a lot more...
- Basically there are thousands of institutions world-wide being mandated to change their internal systems and external interfaces

System Architecture View



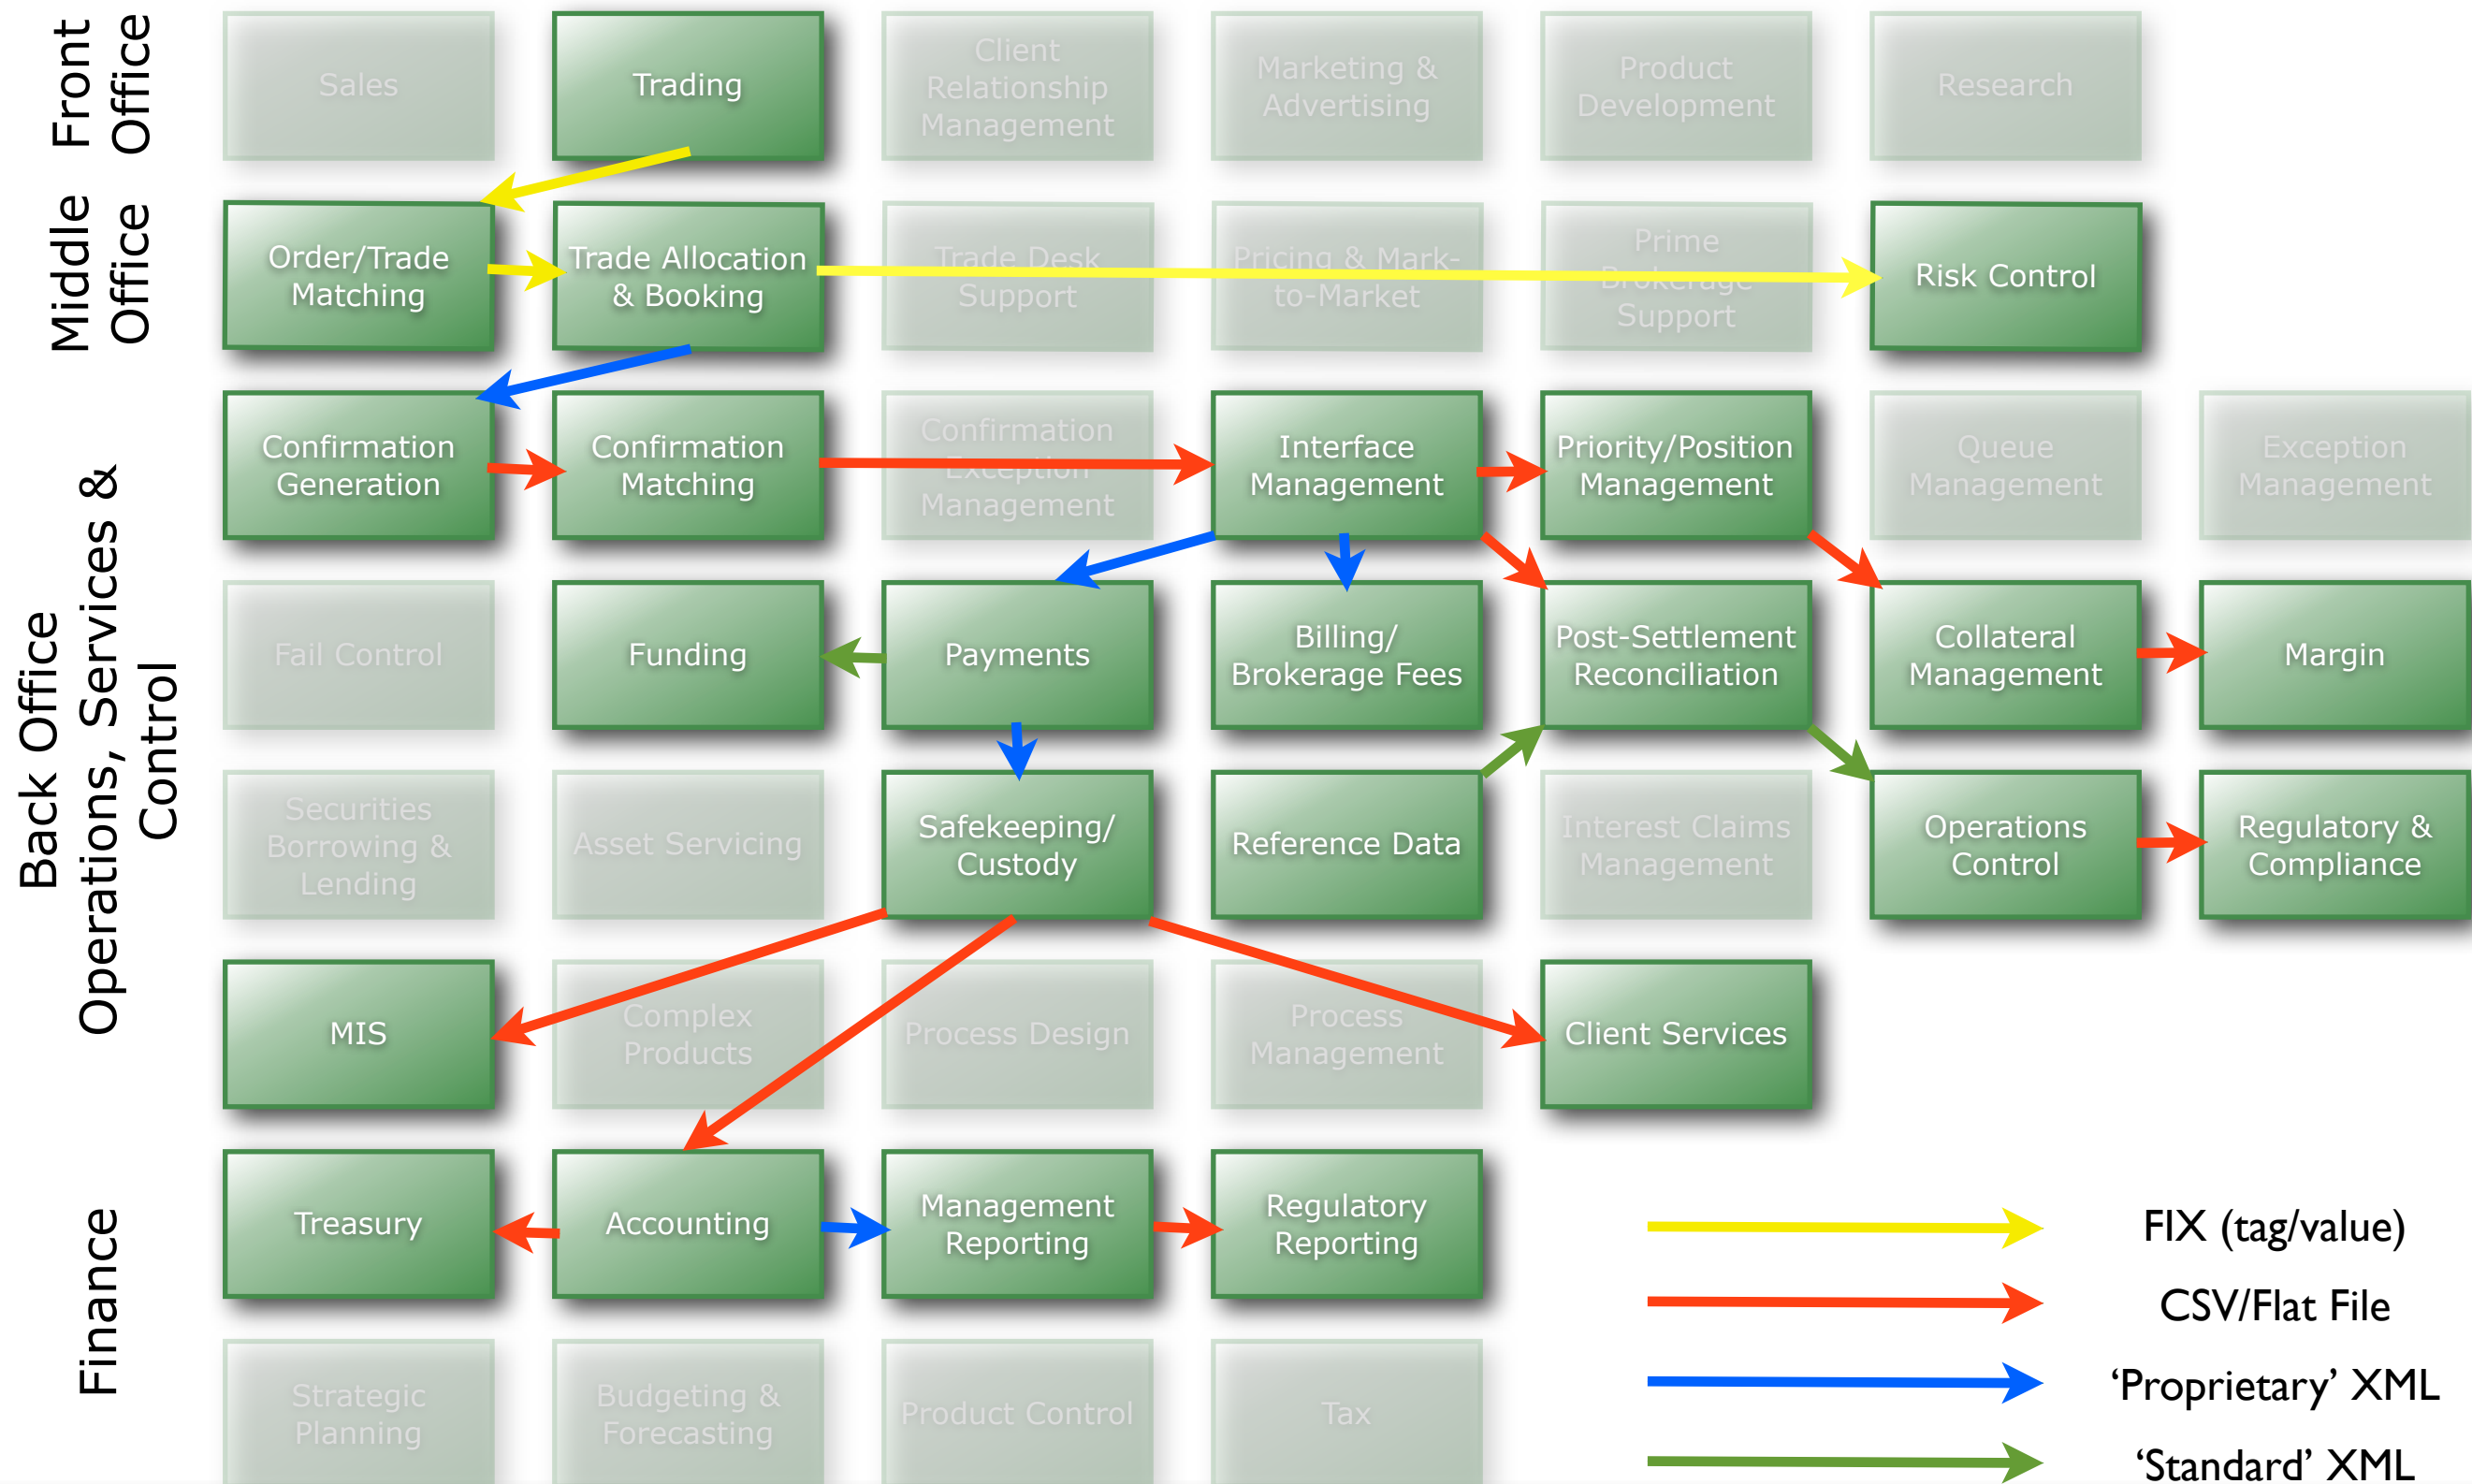
Data everywhere

- Most of the larger banks have literally thousands, many tens of thousands of relational database
- Oracle and IBM spend is usually in 9 (€£\$) figures!
- System age varies between a few months to tens of years
- Data is spread all over the place and in hundreds of different formats, models and schema

Functional Landscape



Just one transaction...



Now add the geography...

US



Europe



Asia Pac



And the rest of the world...

US

Europe

Asia Pac



Add a few standards...

US

Europe

Asia Pac



SWIFT FIN
ISO 20022

Email
HTTPS
mobile

FIX/FIXml/FpML
SWIFT FIN
ISO 20022

FIX/FIXml,
Proprietary

SWIFT FIN,
Proprietary

*ml
XBRL

*ml
XBRL

FIX/FIXml/FpML
Omgeo CTM/OG/OGD
ISO 15022



Corporate
Customers



Private
Customers



Market
Counterparties



Exchanges &
Market Data
Providers



Local Market Agents,
National & International
Payment networks



Central
Banks



Regulatory
Bodies



Clearing & Settlement
Utilities

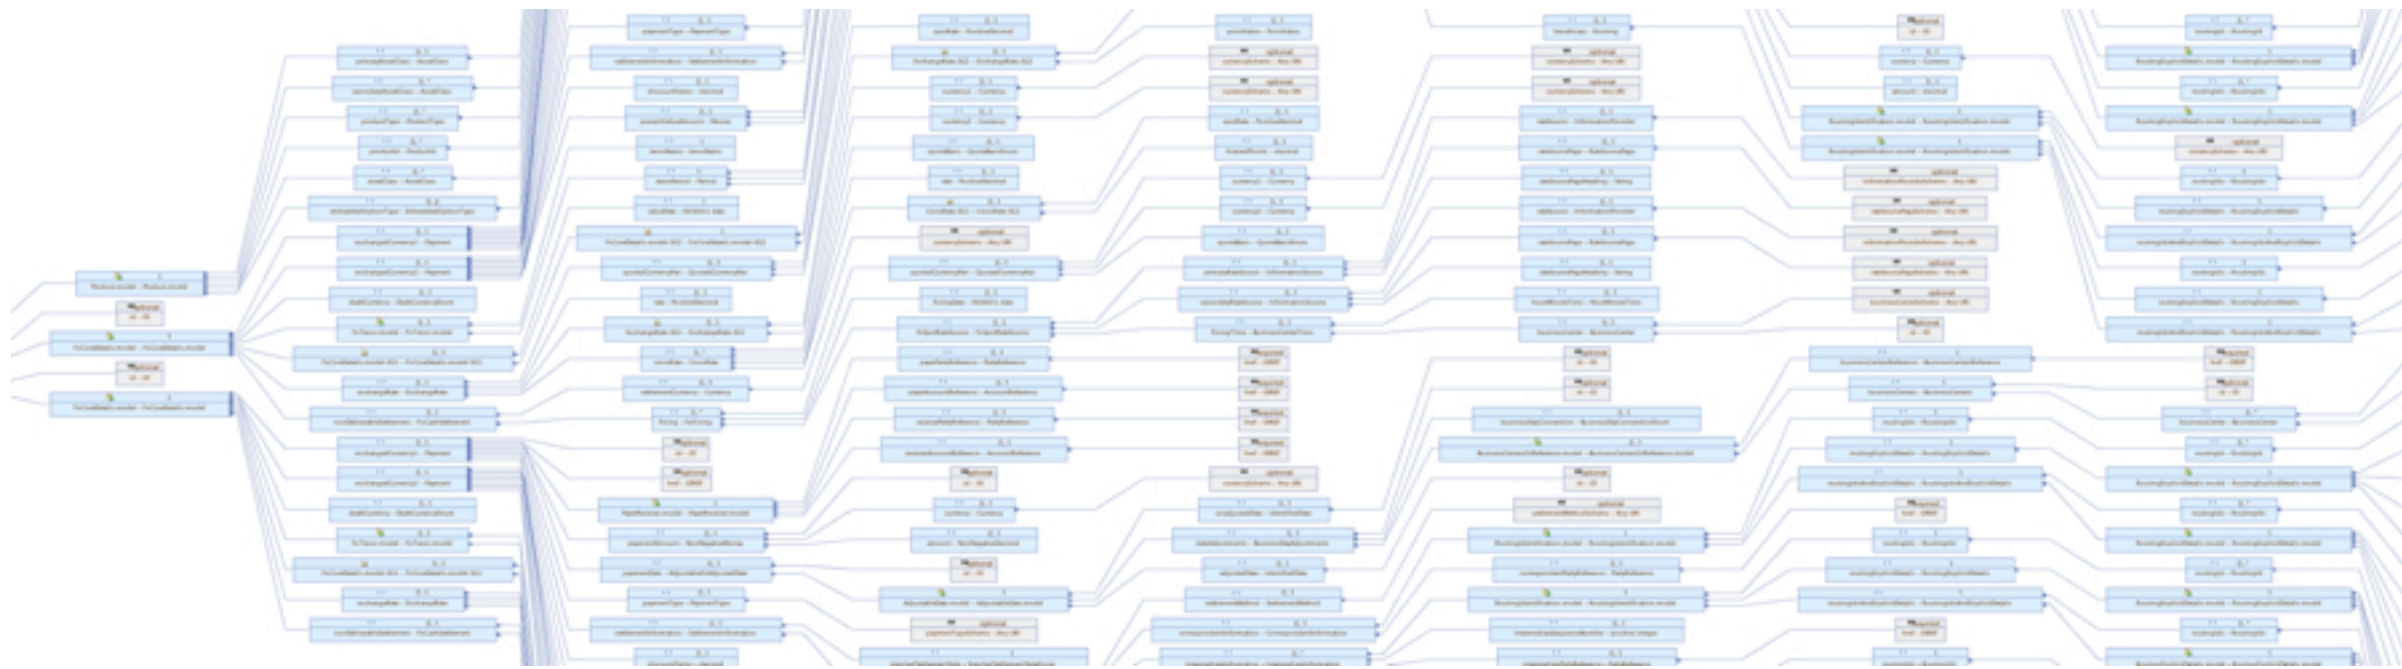
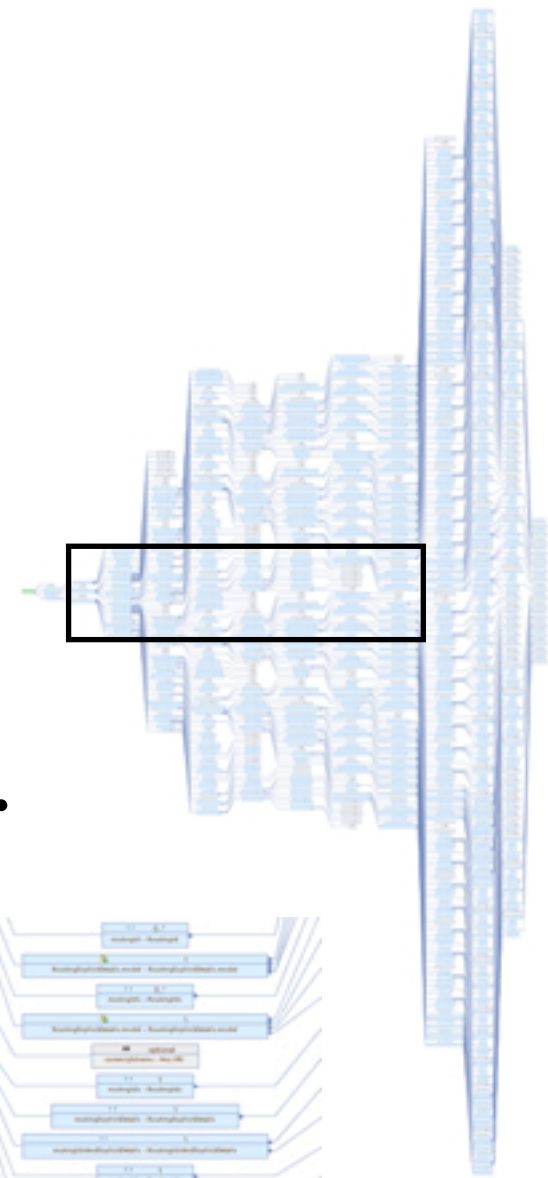
A quick look at FpML

- An FX Swap

- 14 Level of hierarchy
- Over 3,000 elements

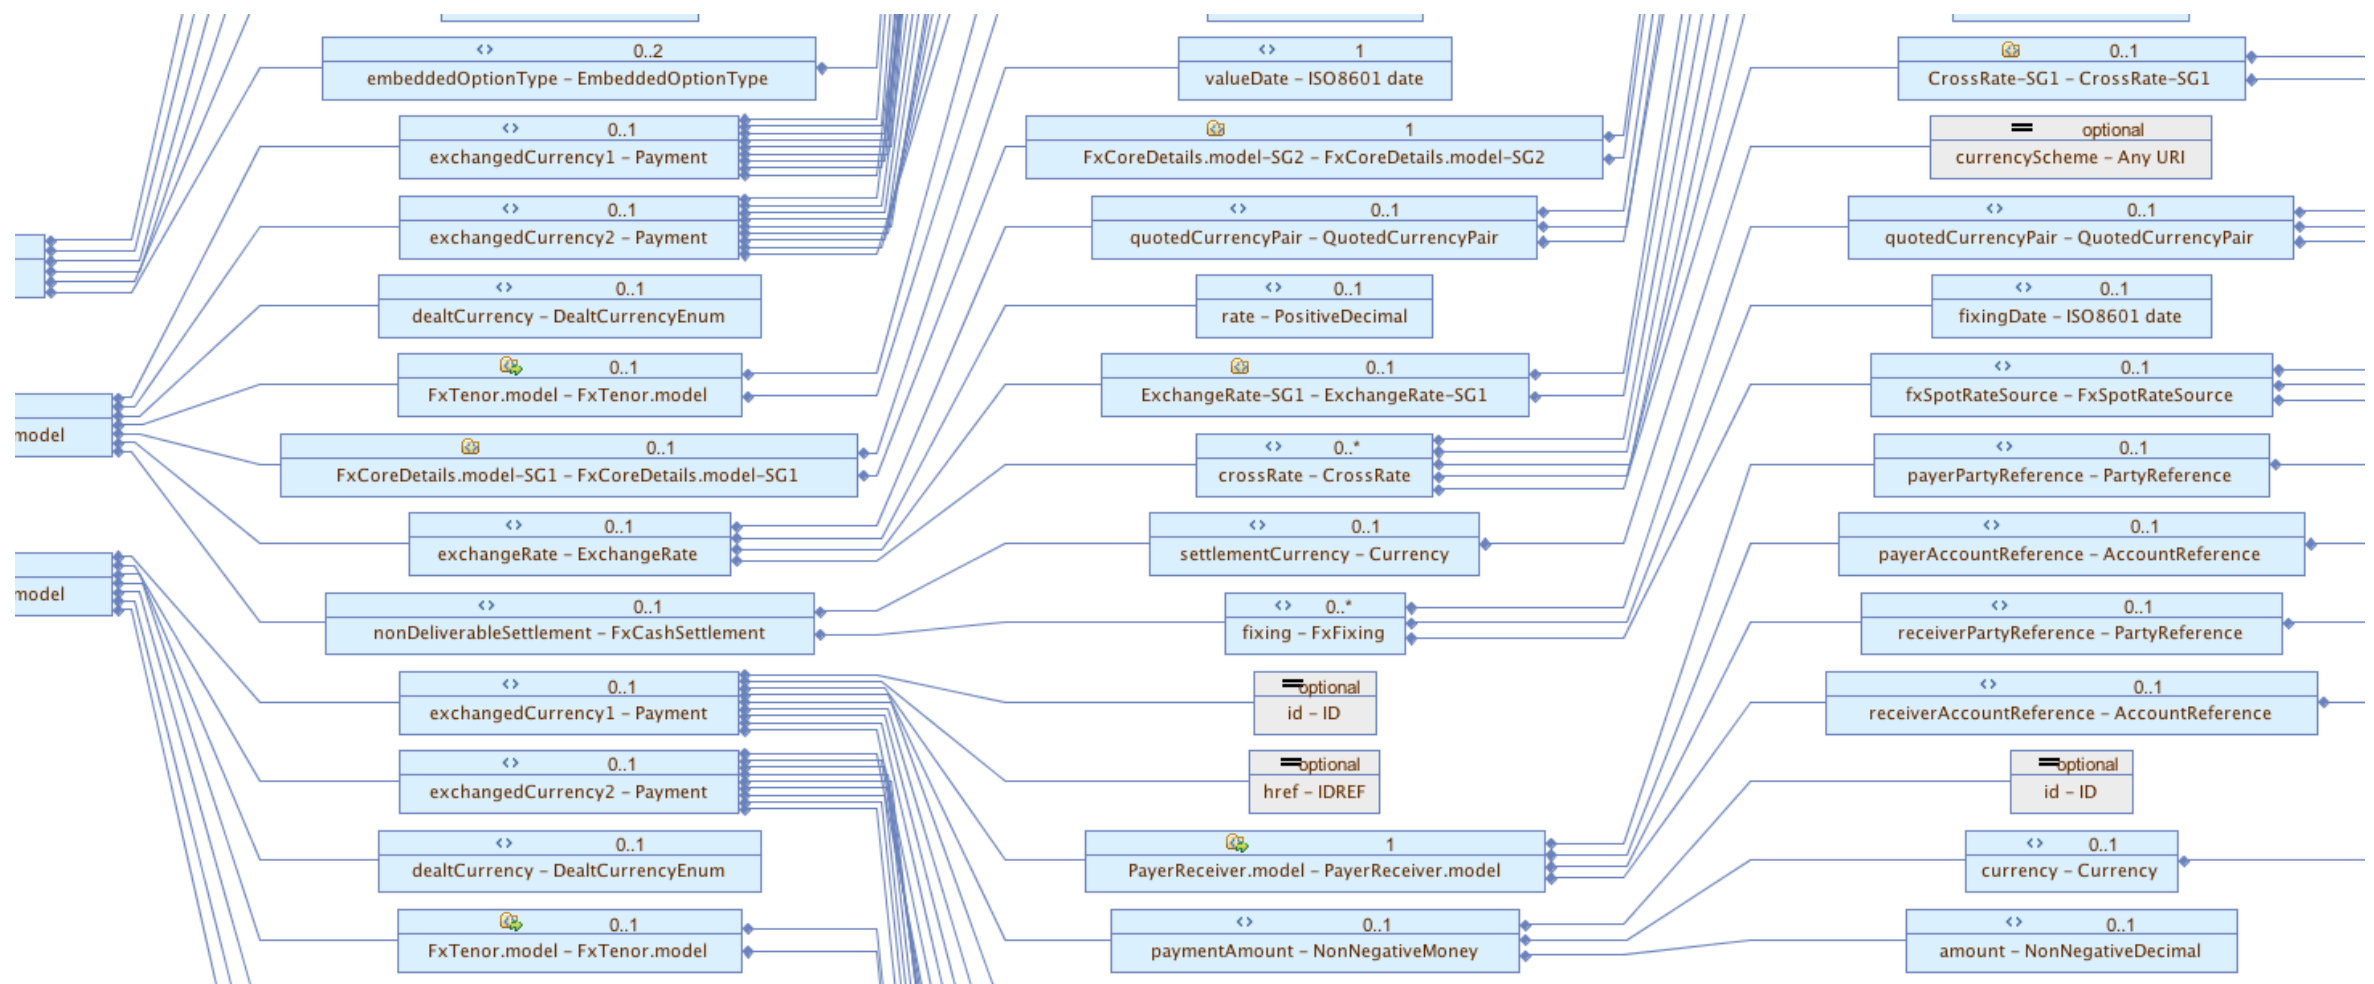
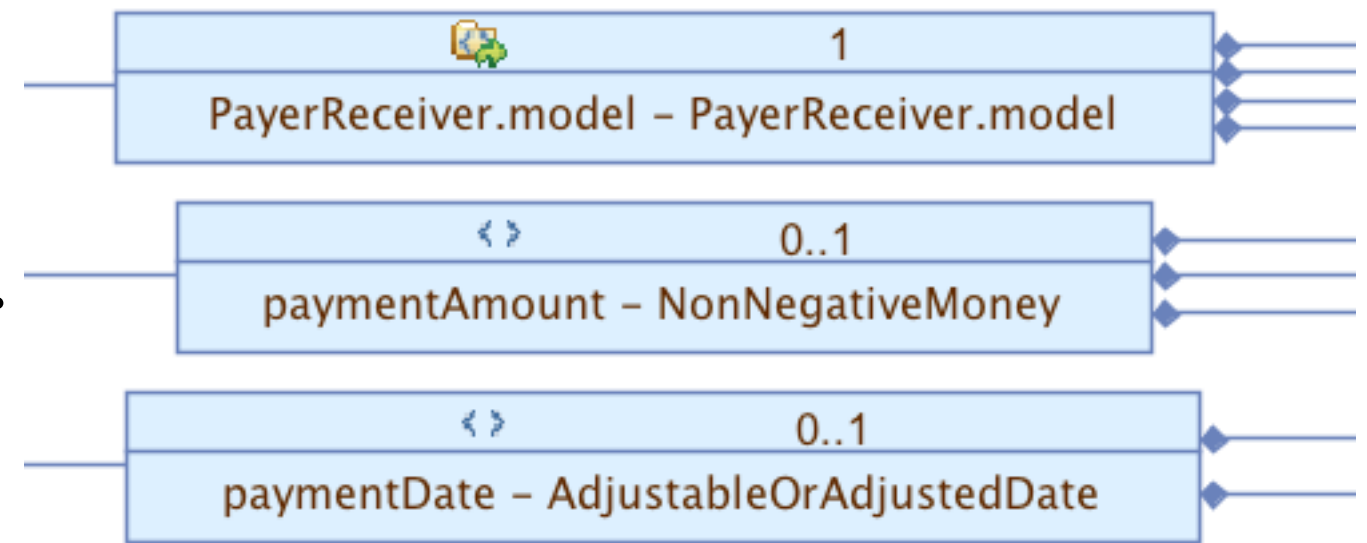
- To the right is the fill message zoomed out

- Below is the part in the box zoomed in a little...



Yes FpML is Complex

- Zooming in a little further...



Not always so bad though

- The schema describes the worse-case but many of the simpler “contracts” (XML instances) are vastly simpler
- The FX swap on the right is pretty much all of the information needed to describe the contract
 - There are no options in this example

```
<trade>
  <tradeHeader>
    ...
    <tradeDate>2002-01-23</tradeDate>
  </tradeHeader>
  <fxSwap>
    <productType>FxSwap</productType>
    <nearLeg>
      <exchangedCurrency1>
        ...
        <paymentAmount>
          <currency>GBP</currency>
          <amount>10000000</amount>
        </paymentAmount>
      </exchangedCurrency1>
      <exchangedCurrency2>
        ...
        <paymentAmount>
          <currency>USD</currency>
          <amount>14800000</amount>
        </paymentAmount>
      </exchangedCurrency2>
      <valueDate>2002-01-25</valueDate>
      <exchangeRate>
        ...
        <rate>1.48</rate>
      </exchangeRate>
    </nearLeg>
    <farLeg>
      ...
    </farLeg>
  </fxSwap>
</trade>
```

To the Database

- If we were just storing FX swaps in a database we'd have no problem
- Twenty years ago this is exactly what we did, we had database tables for each type of trade
 - Another for the currencies, counter parties, nostos, holidays etc. etc.
- But as the derivative market started to mature in the early 2000 we needed better ways to describe the options
- As we continue the relational database ends up being a real problem



ORM - OMG!

- Object Relational Mapping (ORM) is sheer craziness!
- The ORM version of the FpML swap has well over 1,000 tables and a single join is several 'k' in size
- We could create new tables for each contract but that's what we started doing in 2000 and that didn't work
 - Many of these systems are what we have today and this is causing more and more pain
- ORM - Hibernate, JPA etc. was designed for simpler cases



We've got an RDBMS let's use it!

- The usual answer is that we already have Oracle and we pay a lot of money for it so let's use it
- Oracle has an XML type too so what's wrong with that?
- The problem is that Oracle doesn't really want you putting hierarchical data into it's relational database
 - They provide a solution but it's not very performant and very proprietary - so you continue to be locked in
- Could we design a better solution but still using Oracle?
 - BTW - For "Oracle" I also include Sybase, DB2 etc.

Store the FpML in a CLOB

- Back in 2005/6 a few customers started to look at putting their XML into Oracle
 - We looked at native “Oracle XML” and a more generic API
- With this API we parsed the XML, extracted a few key elements and stored those as keys with the XML as the value (in a CLOB)
- And it worked!
- This became the first working version (in financial services production) that I’m aware of

Problem solved?

- Did this solve the problem? - YES
- But why are we using a relational database to store key/value pairs?
 - Worse why are we using a horribly expensive relational database?
- The main reason is that they've already got Oracle and why should they spend more money on something new?
- So we have to wait until we have new issues to solve...

Performance

- Imagine tens of thousands of FpML trades per day, many of them last for years - tens of millions in total
 - Volumes increasing day by day
 - New types of trade every few months
 - New regulations
 - New ways to calculate value and risk
- Needless to say our Oracle DB is starting to become the bottleneck
 - Coherence is one solution we've seen many banks use but it's expensive and "new"



Two new Solutions

- Assuming we're going to move away from the classic RDBMS we now have two new technologies to look at

- NoSQL DBs - Hierarchical data storage



- In-Memory - Relational or Hierarchical but we're interested in Hierarchical here



Here's the flow...

1. Find, gather, read all the trades/messages

- Any ESB/SOA will do - Mule, Fuse, Camel, Spring Integration etc.

2. Extract the information you need for indices

- We use C24's Integration Objects for FpML binding to Java

3. Write the trades/messages to your “database”

- In-memory or NoSQL - GemFire, GigaSpaces, MongoDB etc.

4. Other requirements...

- Ad-hoc queries, rules, CRUD facilities - Combine the above

Enough Slides...

- Time to look at some code and a demo...

It's question time...

- John.Davies@Incept5.com
- Twitter: @jtdavies

