Real-Time Analytics on IBM z Systems

Rajeev V. Kamath (rvkamath@us.ibm.com

Date: Apr. 13, 2016, RTP, NC

Agenda

- Trends in Analytics, Infrastructure Demands and Evolution of the IBM z System

- Predictive Analytics – Improve Quality & Timeliness of Decisions

- Traditional Analytics & Associated Issues

- Best Practice – Take Analytics to the Data; Reduce Data Movement

- The Best Real-Time Analytics Environment (Hybrid Transaction and Analytics Platform – HTAP)

- The Real Time Analytics (RTA) Process: Scoring and Business Rules for Fraud Detection and other use cases

- RTA on IBM z System – Value Proposition

- Real-time analytics use cases
    1. Countering fraud, waste, abuse and financial crimes (Banking example)
    2. Additional in-process payment analytics (Insurance claims example)
    3. Predictive customer intelligence / executing the next best action (Banking/Retail)

Trends in Analytics

- Analytics is becoming pervasive in the enterprise and across industries

- Analytics has become a necessity for not only a competitive edge but also in many situations for survival

- Sophistication and areas on influence of analytics is increasing. Accordingly, data from multiple, disparate sources is being analyzed for better results

- Data volumes are growing rapidly and exponentially

- Data being analyzed includes the traditional structured as well as unstructured data

- Rapid response – near real time, low latency analytics with automation can be a competitive advantage as well as a loss/damage prevention solution. Many new use cases developing and becoming inevitable

- Automated response to the findings of Analytics especially at the operational/transaction level

- Big Data analytics is becoming more common

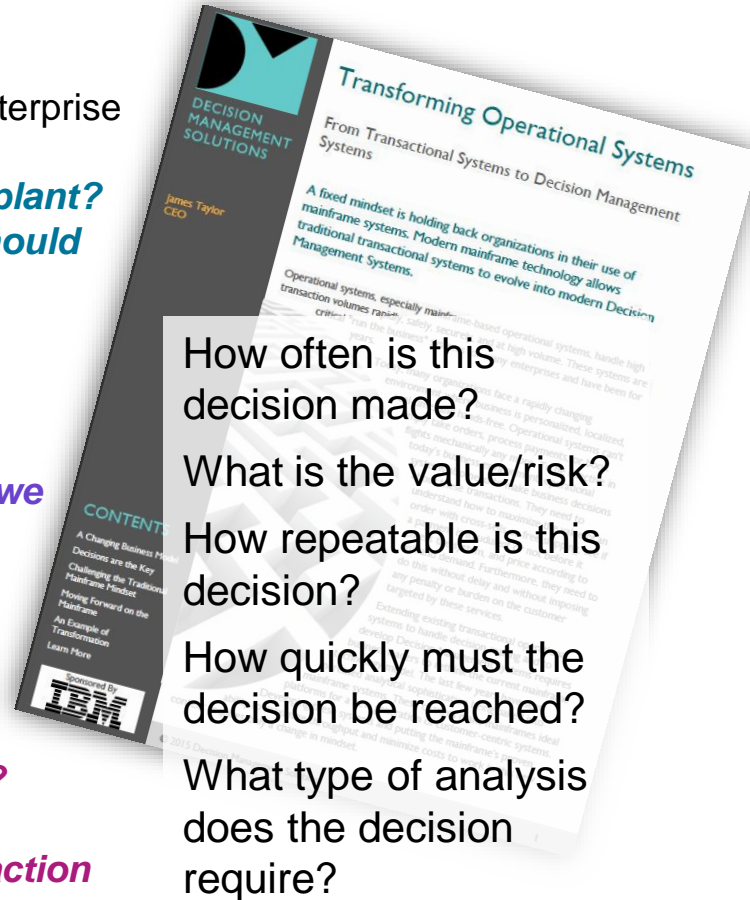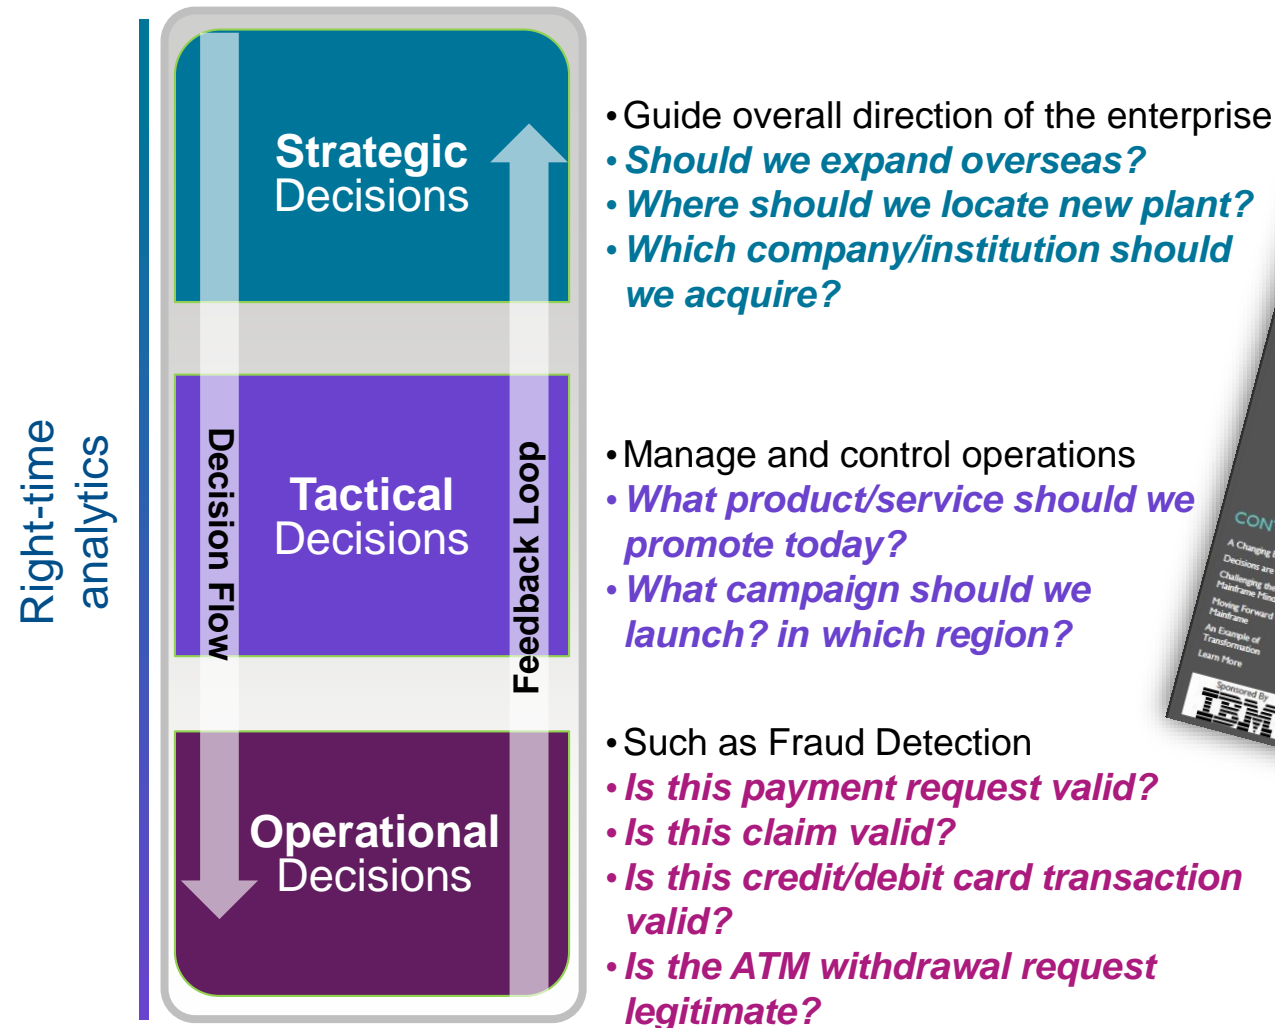# Analytics Trends Are Driving Infrastructure Requirements

- Near continuous availability of data and analytics tools

- Ability to efficiently manage, govern and analyze large amounts disparate data from diverse sources

- High security

- Low latency, near real time, in transaction analytics

- Support for Bigdata Analytics – Hadoop, Spark

- Reduced complexity

- Lower costs

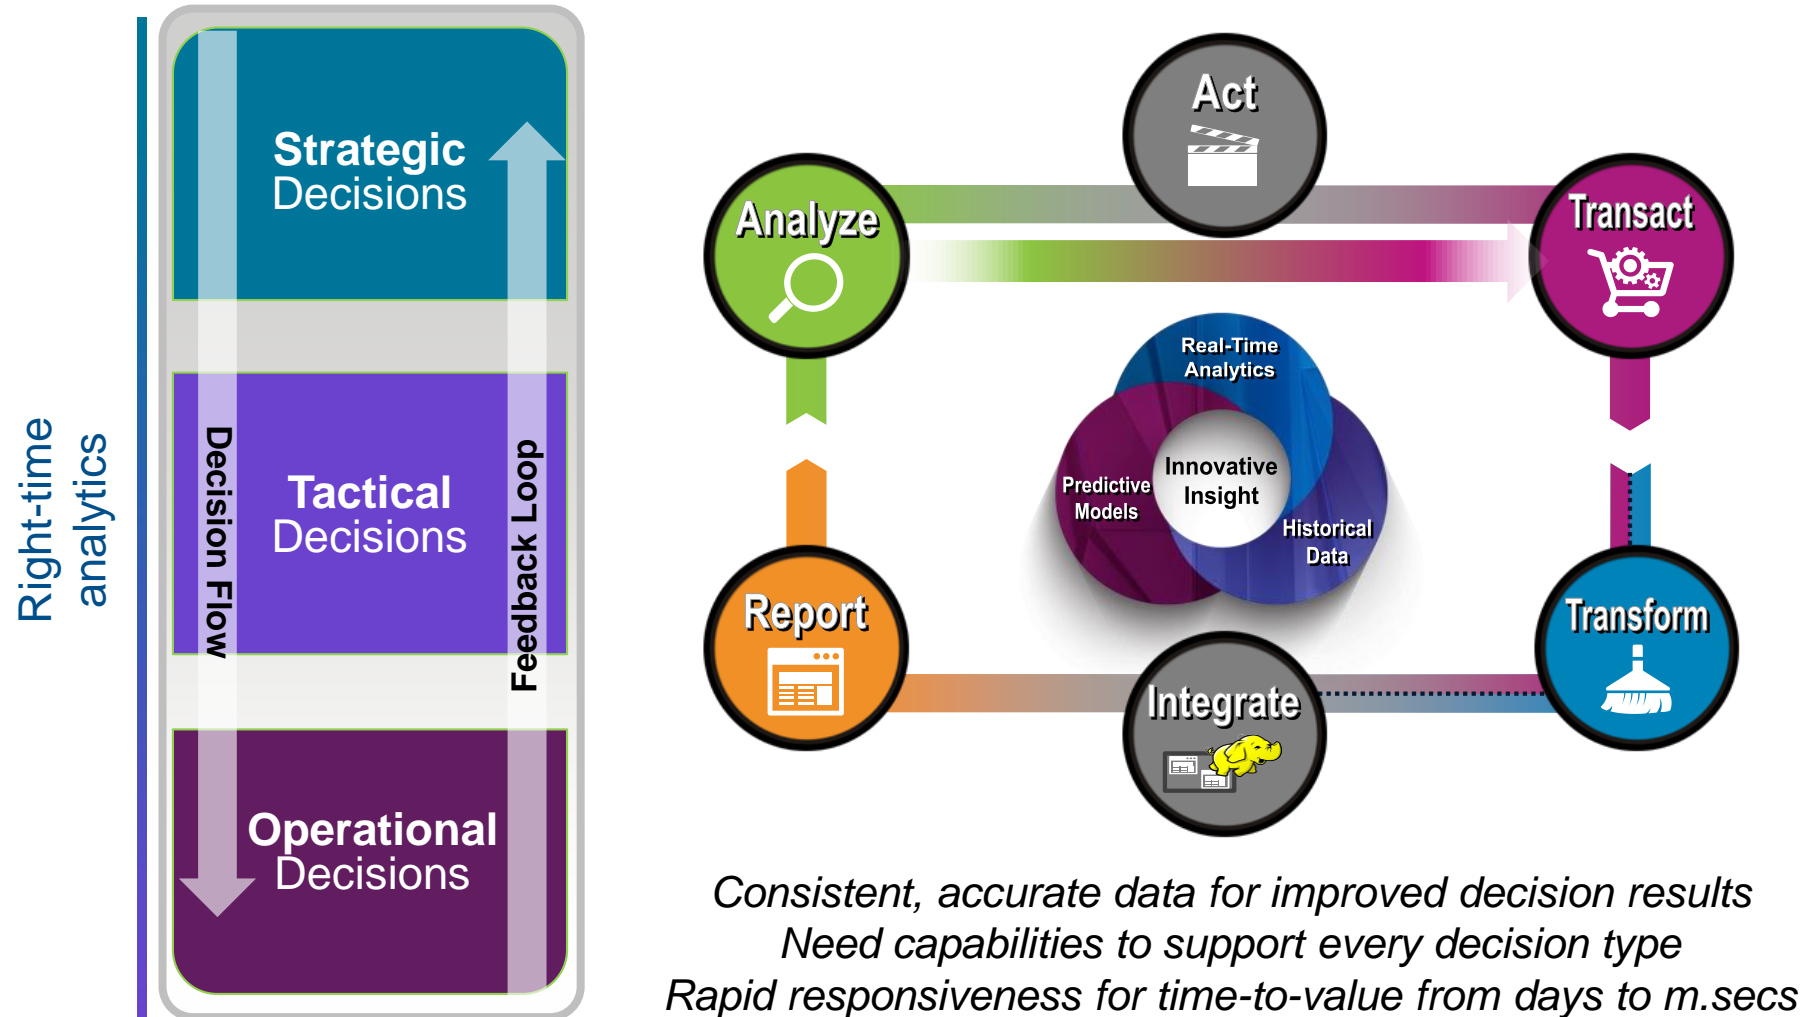# z System Has Evolved In To a Powerful Analytics Platform

- Traditional high qualities of service
  - Near continuous availability – through mulita-level redundancy
  - z System built for Data - Industry leading data management, governance and integrity
  - High security
  - Consolidation capabilities and mixed workload management
  - High scalability
  - Central, simplified system management, simplified backup, restore and DR
  - Higher utilization
  - Better Total Cost of Ownership (TCO)
  - Foundation for cloud, mobility

- Fastest CPU in the industry…with very good floating point performance

- Real Time Analytics – SPSS Scoring Adapter for z/OS, Zementis

- Big Data
  - Integrate with Hadoop analytics
  - Run Spark on z/OS – With data federation from multiple sources, to reduce data movement

Analytics Improve the Quality of Decisions

Types of decisions made in an enterprise.

Right-time analytics

Decision Flow

Feedback Loop

Strategic Decisions

- Guide overall direction of the enterprise
- *Should we expand overseas?*
- *Where should we locate new plant?*
- *Which company/institution should we acquire?*

Tactical Decisions

- Manage and control operations
- *What product/service should we promote today?*
- *What campaign should we launch? in which region?*

Operational Decisions

- Such as Fraud Detection
- *Is this payment request valid?*
- *Is this claim valid?*
- *Is this credit/debit card transaction valid?*
- *Is the ATM withdrawal request legitimate?*

How often is this decision made?

What is the value/risk?

How repeatable is this decision?

How quickly must the decision be reached?

What type of analysis does the decision require?

Transforming Operational Systems
From Transactional Systems to Decision Management Systems

Source: James Taylor, Decision Management Solutions: *Transforming Operational Systems From Transactional Systems to Decision Management Systems*, *2015*

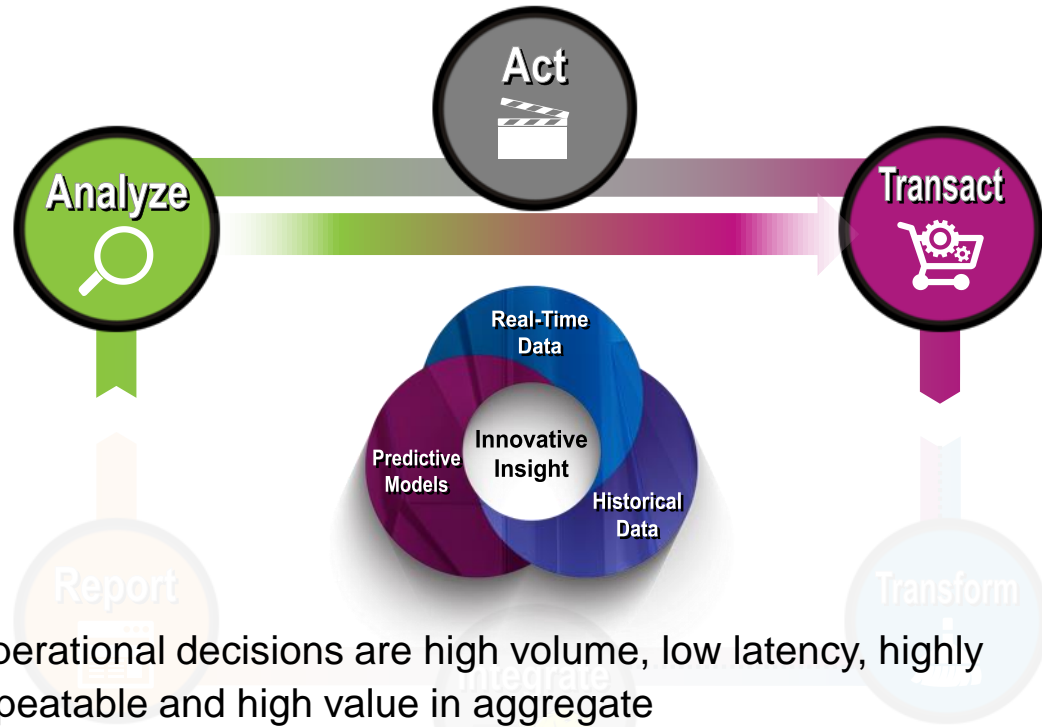# An enterprise analytics architecture is required to transform IT systems into Decision Management Systems

Consistent, accurate data for improved decision results
Need capabilities to support every decision type
Rapid responsiveness for time-to-value from days to m.secs

Today's Focus: Injecting Insight Into Operational Decisions such as Fraud Detection

Right-time analytics

Real-time analytics

Operational decision management is a business discipline, supported by operational and analytics software, that enables organizations to automate, optimize and govern repeatable business decisions to improve the value of customer, partner and internal interactions

Operational Decisions

Analyze

Act

Transact

Real-Time Data

Predictive Models

Innovative Insight

Historical Data

- Operational decisions are high volume, low latency, highly repeatable and high value in aggregate

- They require analytics that are close to the transactional systems so that decisions can be made in **real-time,** and can be automated

- Key considerations: Richness of analysis and ability to maintain SLAs

What is <u>Business Critical</u> Operational Analytics?

Where Right Time = Near Real Time

- An **analytics application that is tightly integrated with OLTP systems** and critical to the optimal running of a business

- Deliver business insight in **real time** or near real time and make decisions in time to **reduce/prevent damage**

- Unavailability of these analytics applications can result in lost business or damage

- Typically support a large, concurrent user population with high volume of requests

Fraud Prevention

Reduce Customer Churn

Cross-sell, up-sell to customers

Real-time Operational Reporting

These applications require an integrated infrastructure with high reliability, availability, scalability and very low latency
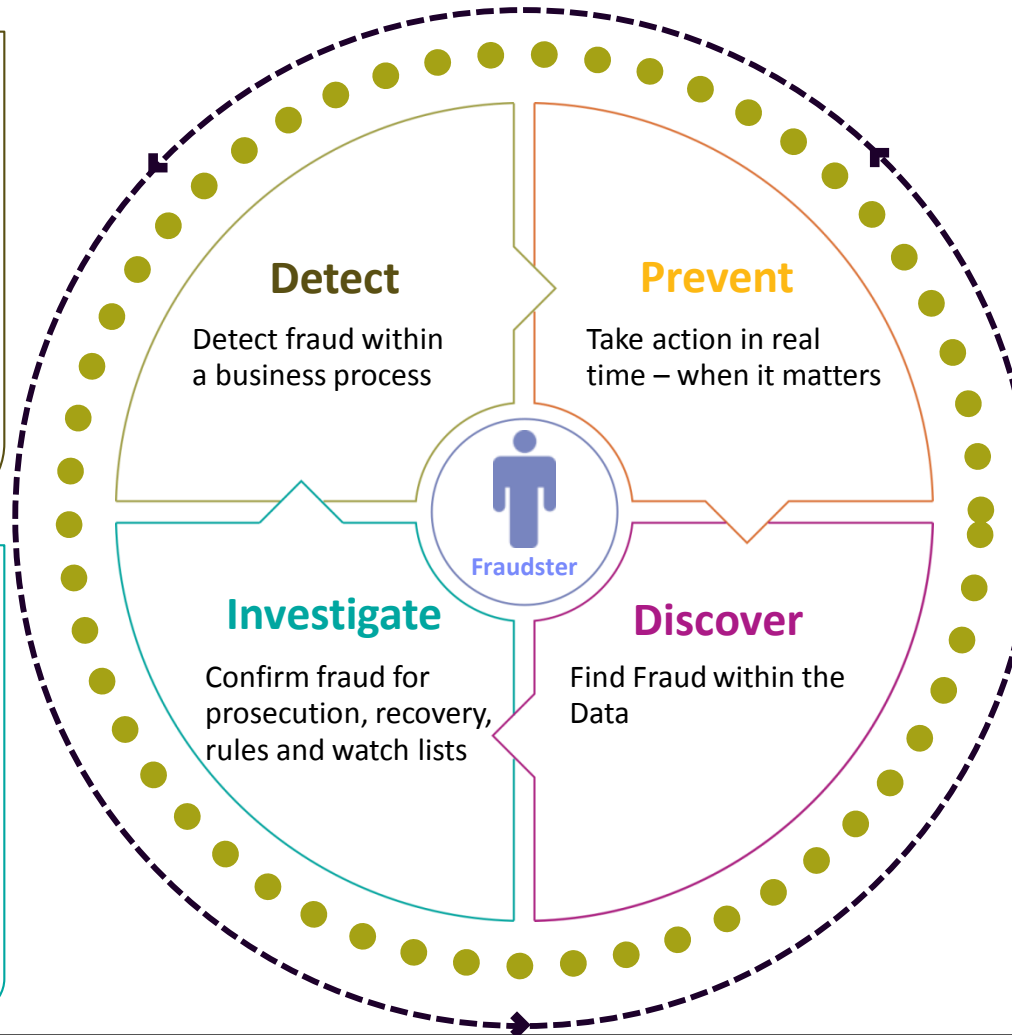
IBM's Fraud Management Point of View

Detect

Detect in real time if a transaction, request, application, document, etc. is potentially fraudulent by applying models and rules in real time

Prevent

Stop processing known fraud, or encourage fraudsters to abandon their objective by showing more is known than they think should be known about their activities and intentions

Investigate

Gather data about fraudsters and/or schemes DETECTED or DISCOVERED; build cases for prosecution, recoveries, or denial of payments. Build watch lists and rules

Discover

Discover fraud by retrospectively reviewing past data and identifying individuals or organizations that may be conducting fraudulent activities

Detect

Detect fraud within a business process

Prevent

Take action in real time – when it matters

Investigate

Confirm fraud for prosecution, recovery, rules and watch lists

Discover

Find Fraud within the Data

Fraudster

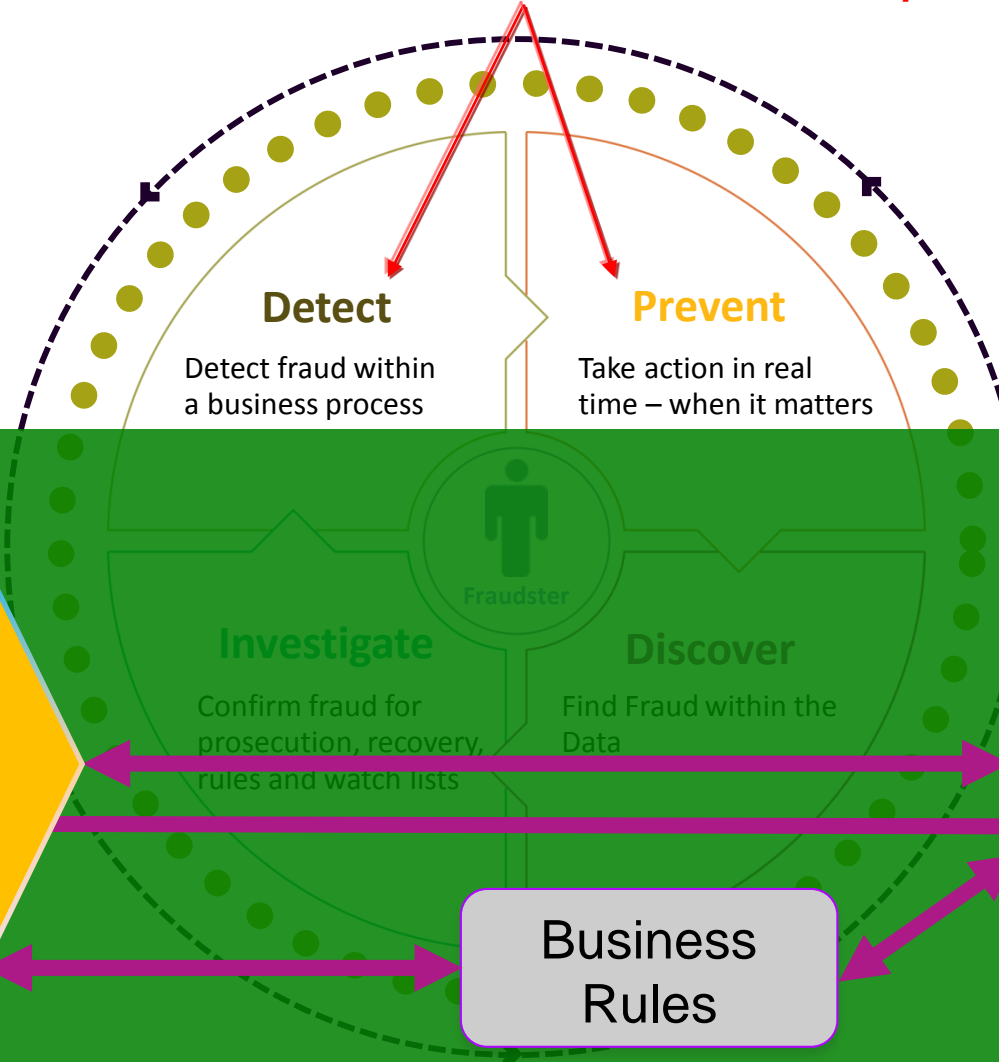# IBM's Fraud Management Point of View

z Systems optimization focus: best results based on location of operational data

Detect

Detect in real time if a transaction, request, application, document, etc. is potentially fraudulent by applying models and rules in real time

Prevent

Stop processing known fraud, or encourage fraudsters to abandon their objective by showing more is known than they think should be known about their activities and intentions

Detect
Detect fraud within a business process

Prevent
Take action in real time – when it matters

Fraudster

Investigate
Confirm fraud for prosecution, recovery, rules and watch lists
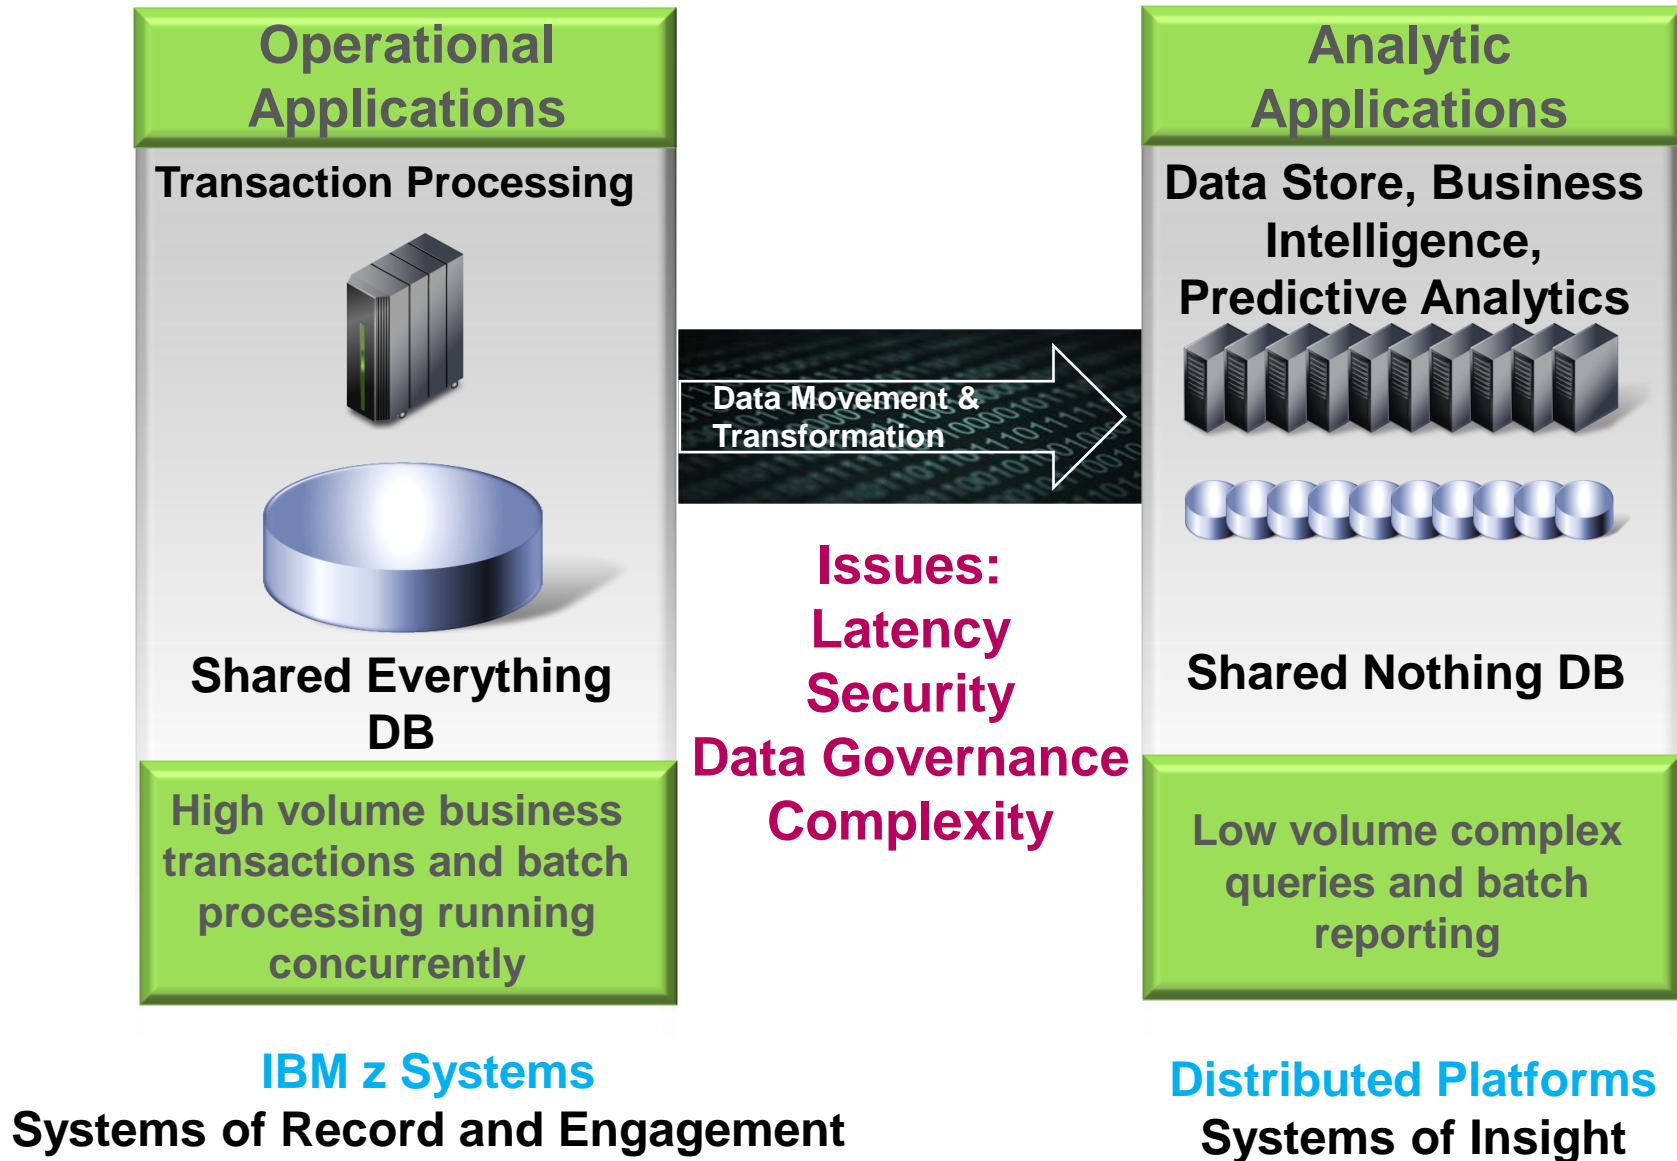
Discover
Find Fraud within the Data

Investigate
Ga... ...abou... fra... ...d/o... sch... ...EC... DIS... ...bl... cas... ...ecut... rec... ...deni... pay... ...ld wat... lis...

Discover
Discover fraud by retrospectively ... ... organizations that may be conducting fraudulent activities

Transaction / Batch Workload

Orchestration

Predictive Analytics

Business Rules

Traditional Approach to Analytics

Operational Applications

Transaction Processing

Shared Everything DB

High volume business transactions and batch processing running concurrently

Data Movement & Transformation

Issues:
Latency
Security
Data Governance
Complexity

Analytic Applications

Data Store, Business Intelligence, Predictive Analytics

Shared Nothing DB

Low volume complex queries and batch reporting

IBM z Systems
Systems of Record and Engagement

Distributed Platforms
Systems of Insight

Actual data flow diagram from an IBM customer engagement (Very common!)
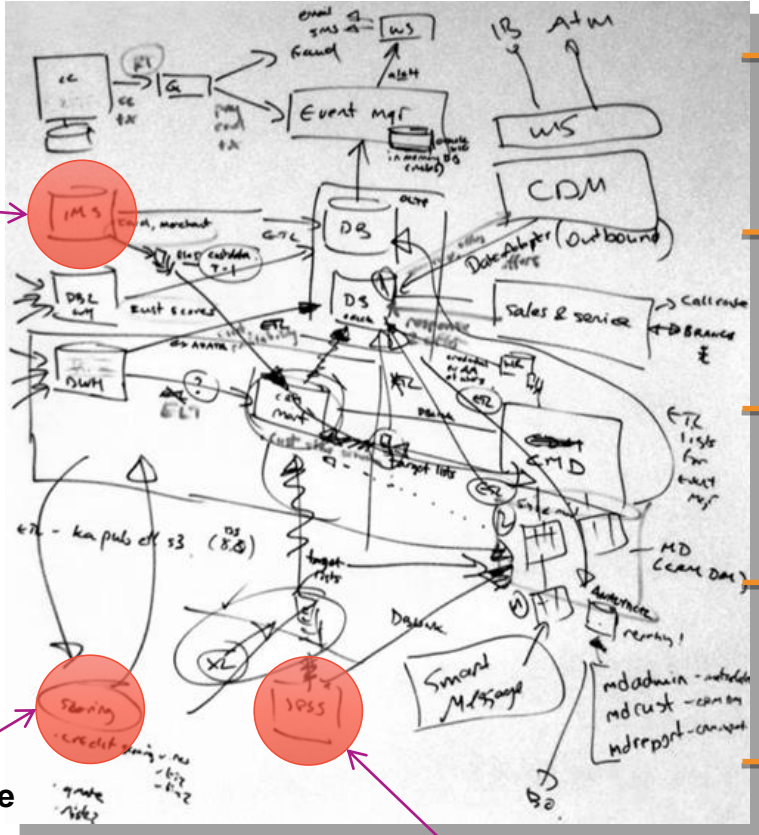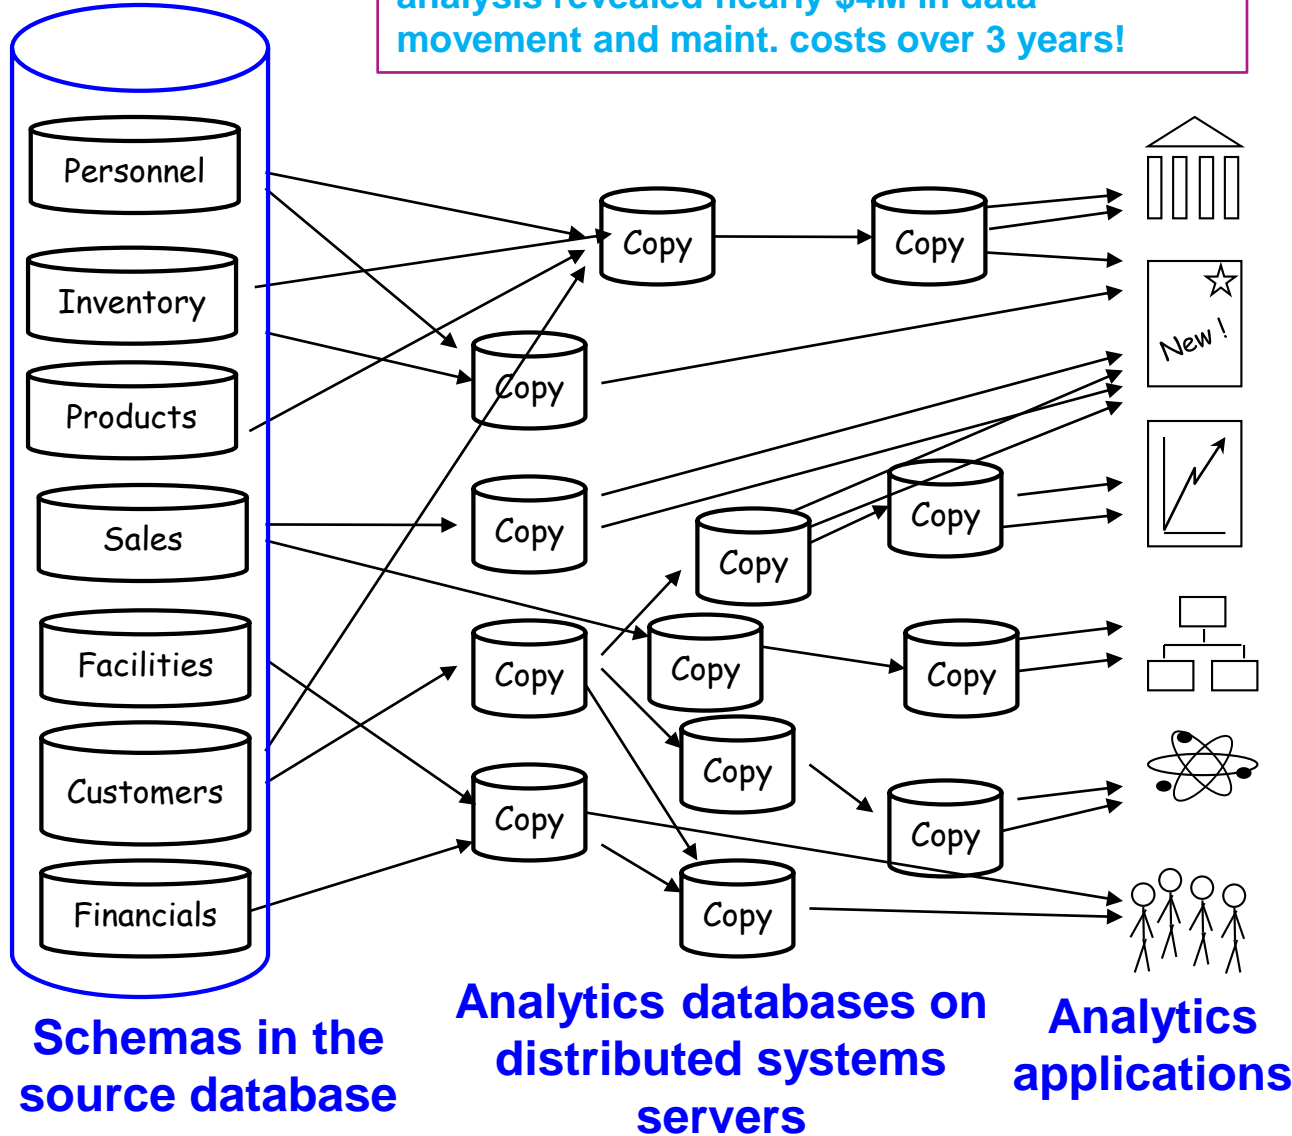

Live Data

Predictive Scoring

Predictive Modeling

Significant complexity
Separate, multiple data warehouses

Analytics latency
Data must be moved, transformed, analyzed & results returned

Lack of synchronization
Multiple copies of data can get out of sync.

Data duplication
Multiple copies of the same data

Excessive costs
Due to data movement, complexity, multiple data warehouse maintenance etc.

Traditional Analytics Typically results in . . .

IBM Customer Engagement Findings:
9 copies of data were being maintained. TCO analysis revealed nearly $4M in data movement and maint. costs over 3 years!

Schemas in the source database

Analytics databases on distributed systems servers

Analytics applications

Issues:

- Data latency: Time between transaction & insight

- Expensive, resource-intensive data movement & transformation processes

- Greater data security exposure

- Data governance issues: Copies of data can become inconsistent – Can users trust the data?

- Data currency challenges: Copies of data can become out-of-date – Users need timely data

- Proliferation of data silos impedes integration, reduces value derived from data assets

The Analyst Community Has Taken Notice!

Gartner®

- "By eliminating analytic latency and data synchronization issues, *hybrid transaction/analytical processing* will enable IT leaders to simplify their information management infrastructure" – HTAP Architecture

- "This architecture will drive the most innovation in real-time analytics over the *next 10 years* via greater situation awareness and improved business agility"

Gartner Research Note G00259033: Gartner 01-2014 Hybrid Transaction Analytical Processing Will Foster Opportunities

The Hybrid Transactional/Analytics Processing (HTAP) Environment on System z

Combined Workloads

Transactional Processing, Traditional Analytics & Business Critical Analytics

OLTP and OLAP On a Single System

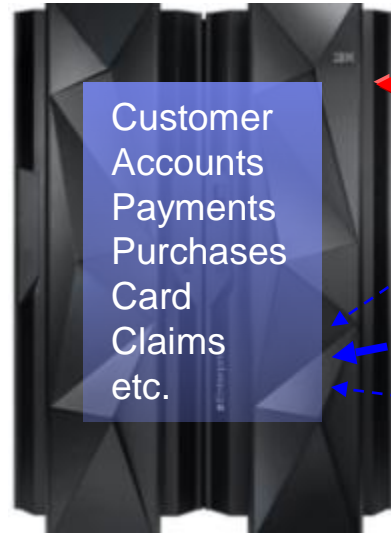**With IBM z Systems HTAP is available now!**

Hybrid DB
Reduced Latency. Greater Security.
Improved Data Governance. Reduced Complexity.

High volume business transactions, real time scoring and batch reporting running concurrently with complex analytics

The Better Infrastructure for Real Time Analytics!

The Better approach!

Operational Systems of Record

Process Orchestration

Predictive Models

Customer
Accounts
Payments
Purchases
Card
Claims
etc.

network

network

network

Data Accessed

Data Moved

Data Accessed

Business Rules

Move the data to the analytics

Place the analytics near data

Operational Systems of Insight

Customer Accounts
Payments
Purchases
Card Transactions
Claims
etc.

Process Orchestration
Predictive Models
Real Time Scoring
Business Rules

Considerations:

Can performance / throughput SLAs tolerate data movement and network traffic?

Can models integrate large volumes of historical data with incoming transactions to deliver the most accurate outcomes?

Can security for sensitive data be maintained across multiple zones?

Can audit trails be maintained to satisfy regulations?

Can availability and BC/DR objectives be met?

Can 100% of transactions be richly analyzed without user impact?

- **Low latency, Real Time**
- **Simplicity**
- **Fewer data copies (Less ETL, Data Governance and Security Issues)**

18

Real Time Analytics – Two Major Objectives

- Detect an undesirable outcome or risk <u>in time to prevent it</u> and/or
- Recognize an opportunity <u>in time to take advantage of it</u>

- Analytics can uncover a threat, risk or an opportunity. In many cases, the time is of the essence

  - *Identify a threat, risk or an opportunity and act on it before the damage is done, or the opportunity is lost!*

  - *Some Examples:*
    - *Banks – Fraudulent credit card or ATM Transaction*
    - *Insurance – Fraudulent claims*
    - *Retail – Opportunity to Cross-Sell Up-Sell*
    - *Telco – Customer Churn*
    - *Government – Tax fraud, Payment fraud*

Real-Time — THE TIME IS NOW

Best Practice - Take Analytics to the Data

Moving data between platforms just for Analytics using network resources and the ETL process, is unproductive, inefficient and expensive!

IBM Customer engagements have shown 20% or more z Systems CPU used just for ETL!

Real-time analytics can be used to maximum benefit when it's executed on the platform where transactions are processed, data is created and stored

- In situations where the data to be analyzed "lives" on z Systems, that means performing real-time analytics on z Systems

- In particular, analytics on z Systems makes in-transaction invocation of predictive models feasible for high-volume, mainframe-based application environments such as CICS-DB2

- Analytics on z Systems also reduces latency and enhances security for the data being analyzed, when the source data is z-based

    - Data movement and replication costs can be lowered versus sending data off-platform for analysis
    - Fewer copies of data improve governance and security

SPSS – Predictive Analytics For Fraud Detection

- Data mining capabilities of SPSS can spot patterns in historical data that humans can not discern
  - When it's not clear as to what question to ask
  - SPSS highlights the areas of focus
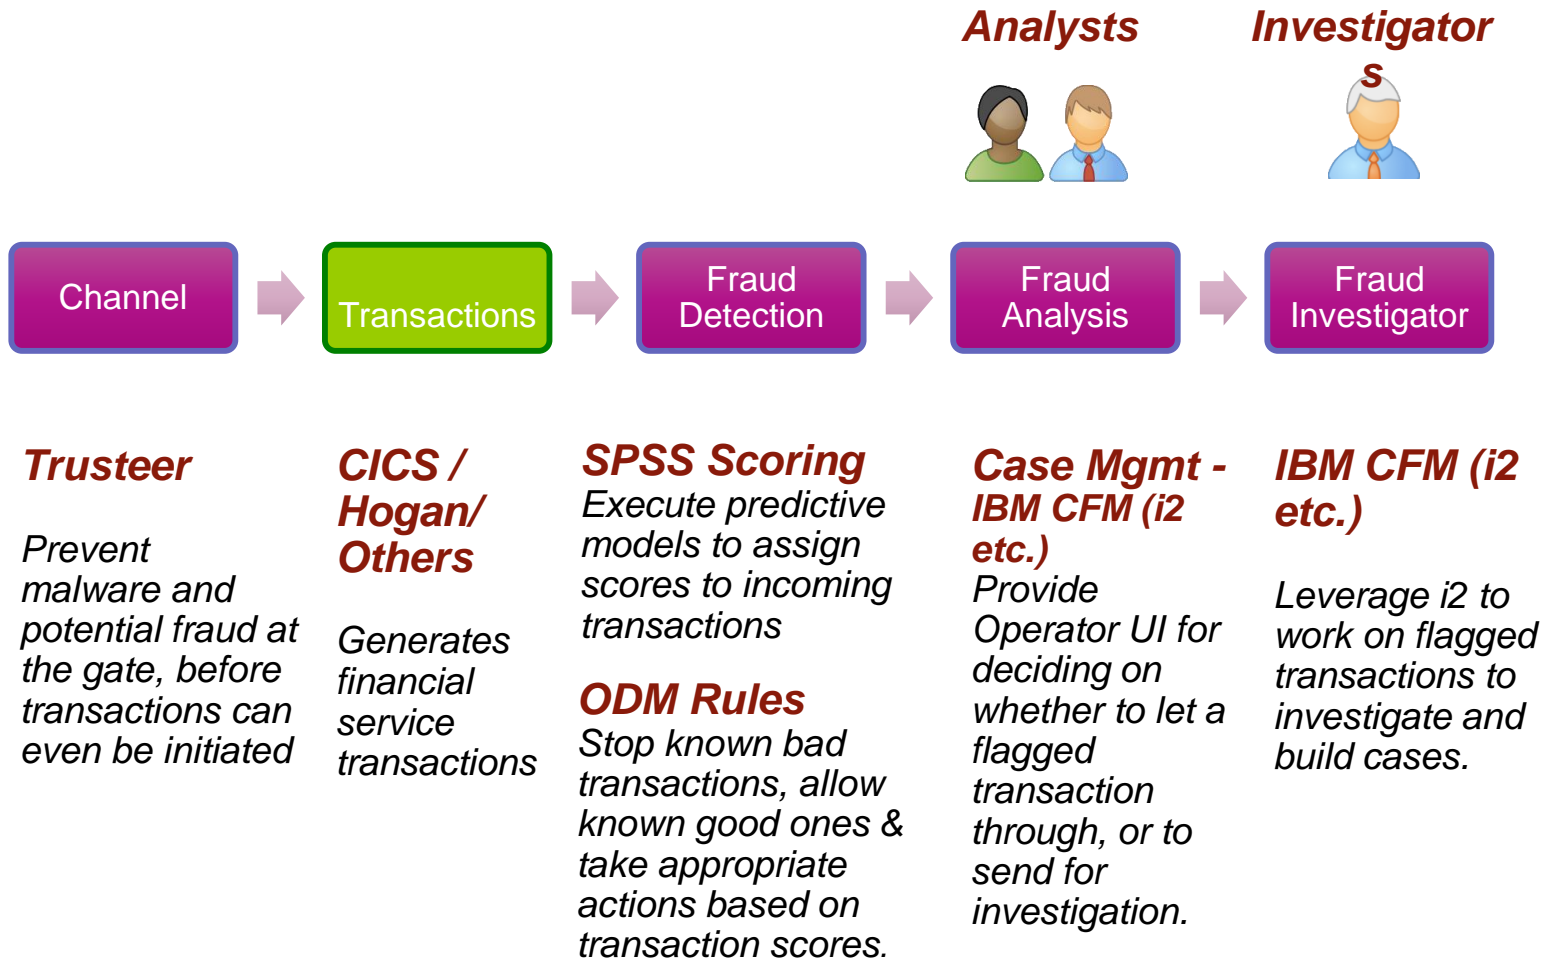 - Identify risks, opportunities and estimate outcomes
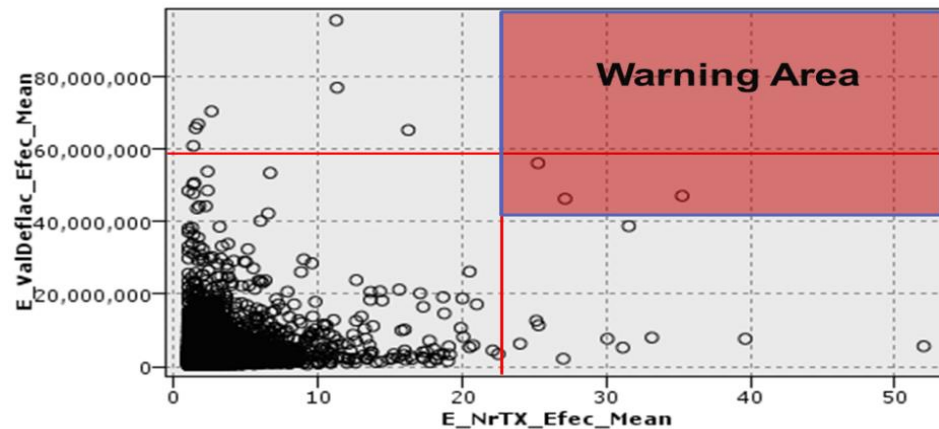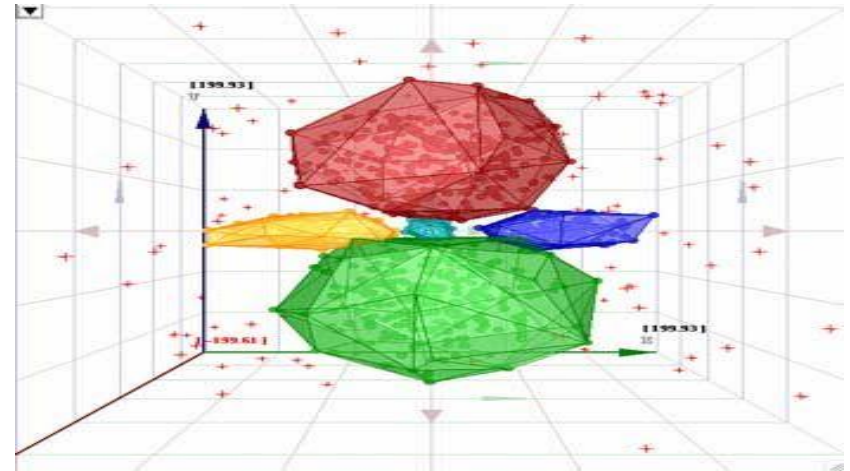
- SPSS can use the detected pattern to develop predictive models that can determine a 'score' for every transaction
  - 'Score' is a mathematical metric that can identify propensity or probability of the transaction being fraudulent
  - Has built in algorithms such as regression, time series, decision trees, k-means, neural networks etc.
  - Easy to use, drag and drop design

2

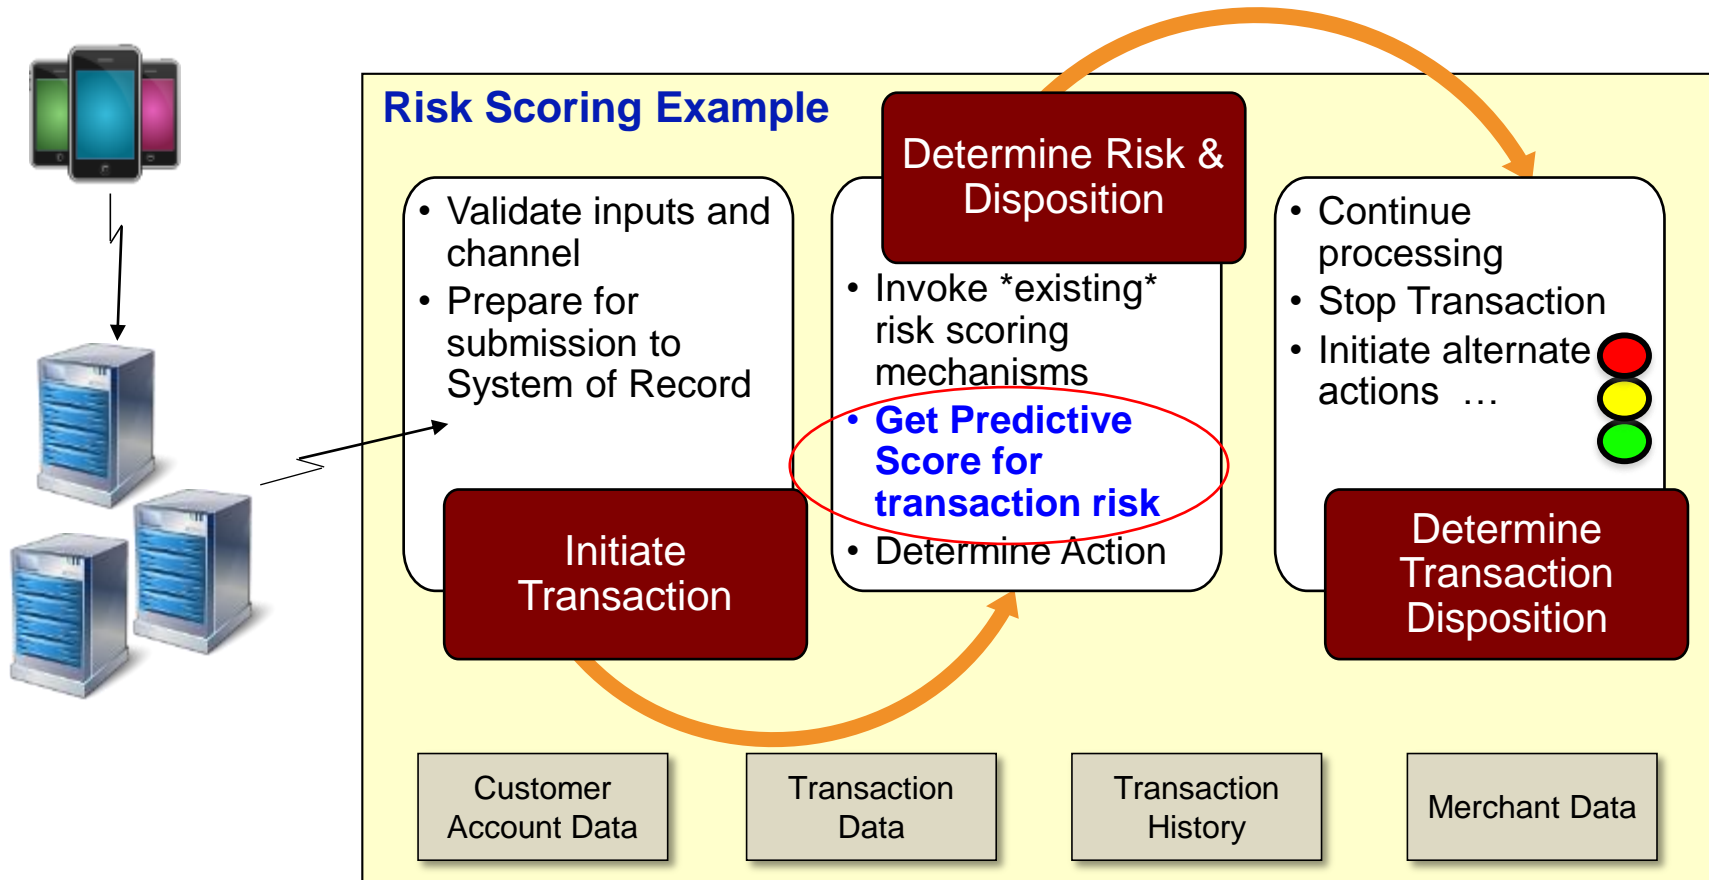# SPSS Scoring Adapter on z Systems for 'In-Transaction', Real Time Scoring

- A 'Score' is a result of mathematical and statistical calculations and is an indicator of the characteristics of a transaction such as a propensity or a probability of an occurrence or an outcome etc. The Score can be used to determine how best to process the transaction!

- The *IBM SPSS Modeler with Scoring Adapter for zEnterprise* facilitates the publishing of an SPSS predictive model to DB2 for z/OS in the form of a user-defined function (UDF)

  – A UDF is essentially a data processing program that can be invoked in an SQL statement – typically a SELECT. It can be included in the transaction processing workflow for execution in near real time!

- By pushing the predictive model's functionality into **on-platform** **data layer** of an application, the SPSS Scoring Adapter provides high-performance, highly scalable, real time <u>in-transaction</u> scoring

  – A transaction is assigned a score *while it is still being processed*
  – Scales to thousands of transactions per second
  – Adds only a few milliseconds to a transaction's processing time
  – Allows an action to be taken *before the transaction processing is completed!*
  – *Instrumental in avoiding the "Pay and Chase' approach!*

IBM's Counter Fraud Solution: High Level Flow

Clients and .. "non-clients"

Analysts

Investigators

Channel → Transactions → Fraud Detection → Fraud Analysis → Fraud Investigator

Trusteer

Prevent malware and potential fraud at the gate, before transactions can even be initiated

CICS / Hogan/ Others

Generates financial service transactions

SPSS Scoring
Execute predictive models to assign scores to incoming transactions

ODM Rules
Stop known bad transactions, allow known good ones & take appropriate actions based on transaction scores.

Case Mgmt - IBM CFM (i2 etc.)
Provide Operator UI for deciding on whether to let a flagged transaction through, or to send for investigation.

IBM CFM (i2 etc.)

Leverage i2 to work on flagged transactions to investigate and build cases.

In-transaction, synchronous predictive analytics

Risk Scoring Example

Initiate Transaction

- Validate inputs and channel
- Prepare for submission to System of Record

Determine Risk & Disposition

- Invoke *existing* risk scoring mechanisms
- **Get Predictive Score for transaction risk**
- Determine Action

Determine Transaction Disposition

- Continue processing
- Stop Transaction
- Initiate alternate actions …

Customer Account Data	Transaction Data	Transaction History	Merchant Data

In-transaction predictive analytics can address issues such as data latency, timely execution, tight SLAs, governance of sensitive data, access to both transaction and external data, skills gaps.

Integration of Real-Time Scoring For Fraud Detection into Transaction Processing Workflow

1. Create a stream to build a model.

2. Execute the model in SPSS to produce a model apply node.

3. Evaluate and test the model with test data in a separate test partition

4. Publish the model to DB2 subsystem (Nugget!)

5. Add model execution to generate a score for every transaction Via DB2 UDF invocation

6. To rebuild the model, repeat steps 1..4 above.
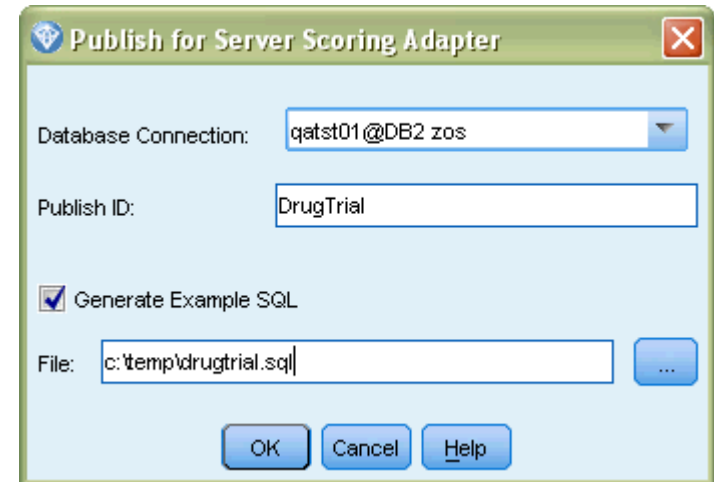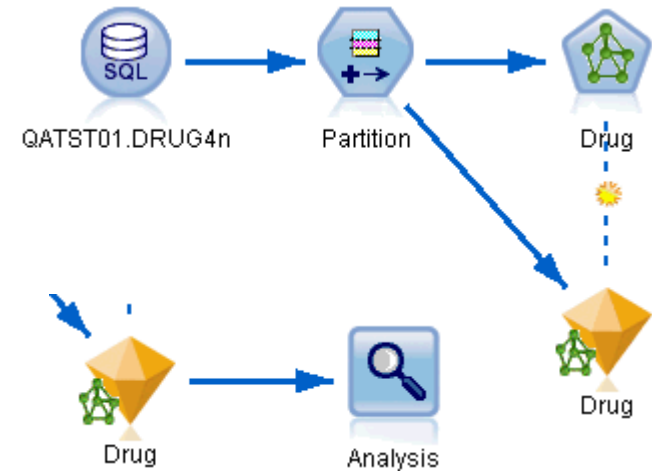
Fraud Models for Insider and Online Fraud

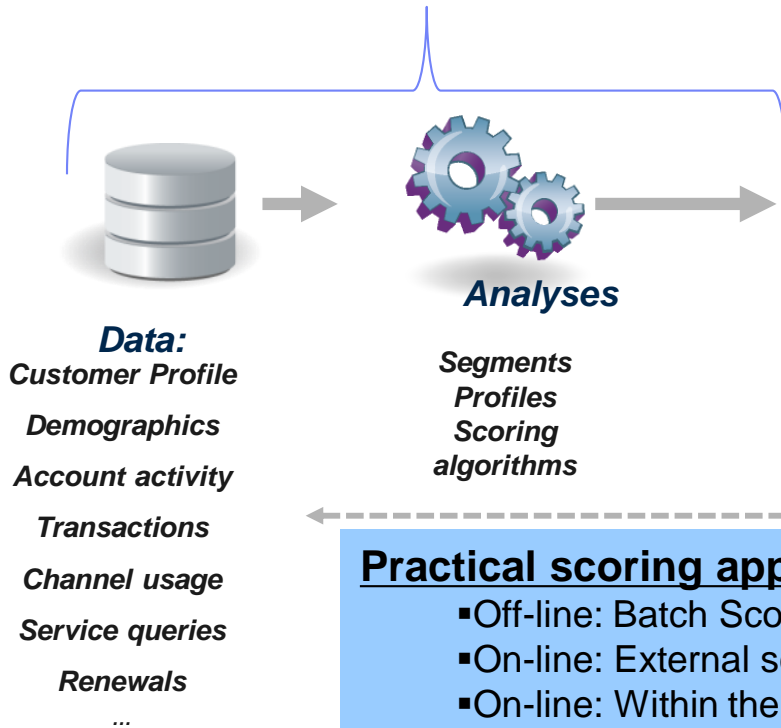## SPSS Modeler stream screenshots for different analytic models
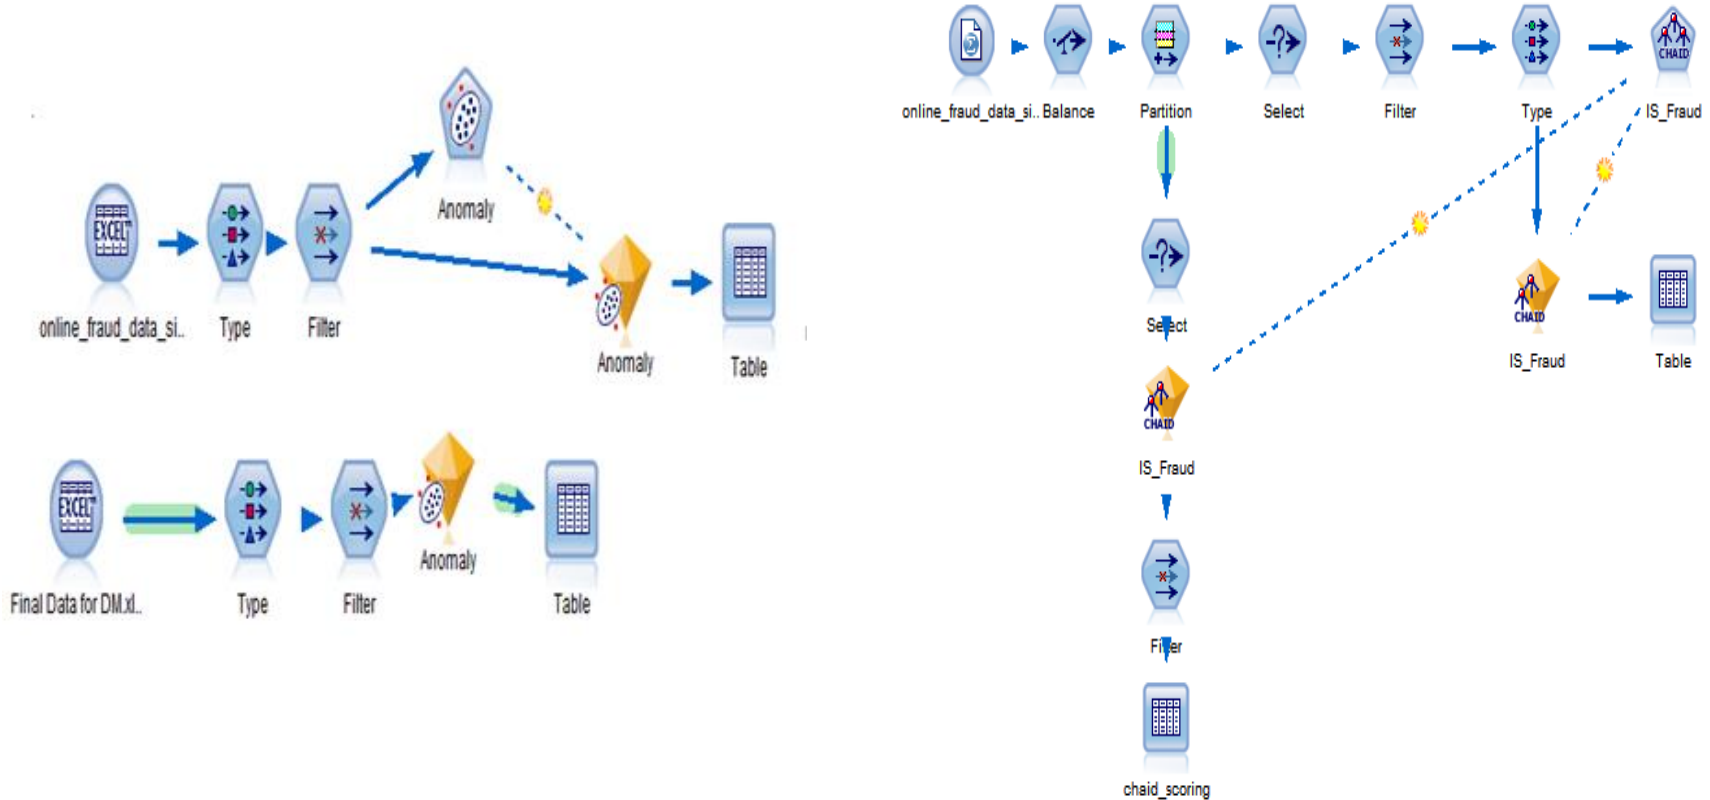
Analyze transactions in **real time**, as part of **transaction processing workflow** for potential fraud using the CHAID decision tree.
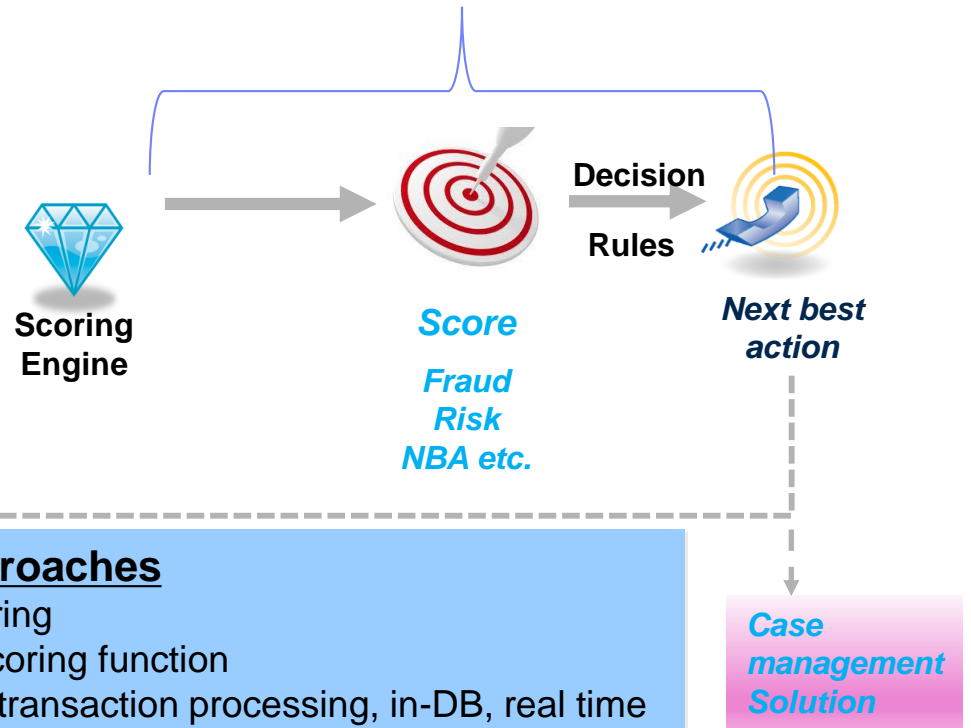
Predictive Scoring Process

Step 1 – Build the predictive model

Step 2 – Execute the predictive model

Identify patterns in historical data and build and test predictive models and a score

Use the predictive models to score every transaction and decide it's disposition

Analyses

Scoring Engine

Decision

Rules

Score

Next best action

Data:
Customer Profile

Demographics

Account activity

Transactions

Channel usage

Service queries

Renewals

...

**Segments
Profiles
Scoring
algorithms**

*Fraud
Risk
NBA etc.*

Practical scoring approaches
- Off-line: Batch Scoring
- On-line: External scoring function
- On-line: Within the transaction processing, in-DB, real time

Case management Solution

Real Time Analytics – The Workflow
Predictive threat & fraud analytics

Transactions

Context Data:
- *Current transaction data*
- *Location, frequency, timing etc.*

Build Model

Predictive Model/ Scoring

SCORE

Rules Engine

Reduce False Positives

Historical Data:
- *Customer Profile*
- *Account Information*
- *Demographics*
- *Activity history*
- *External data*

Decision

Suspect

Refer for investigation

Standard processing

Fast Processing

Legitimate

Scoring Workflow in Traditional (Off z Systems) Analytics

Batch Scoring

User submits transaction

Business Application / OLTP

Historical Data store

Scoring Application

ETL

R-T, min, hr, wk, mth

Copy

Delivers static score

(after the fact, 'pay & chase')

Application

DB

Distributed Severs

z System

Batch Score Copy (nightly, weekly, monthly)

Real-time Scoring

'Real-Time' Web Services Call/Response

User submits transaction

Business Application / OLTP

Historical Data store

Scoring Application

ETL

R-T, min, hr, wk, mth

Copy

Delivers a dynamic score

Application

DB

z System

Distributed Severs

- *Latency & Complexity*
- **Data Governance, Currency Synch and security Issues**

The Impact of (Off z System) Real-time Scoring on IT Infrastructure

Multiple infrastructures required

Web services call

Business Application / OLTP

Historical Data Store

Scoring Application

Customer interaction

Delivers a dynamic score

Application DB

ETL

R-T, min, hr, wk, mth

Copy

High Networking Traffic

Web services call

Difficult to deliver SLA's
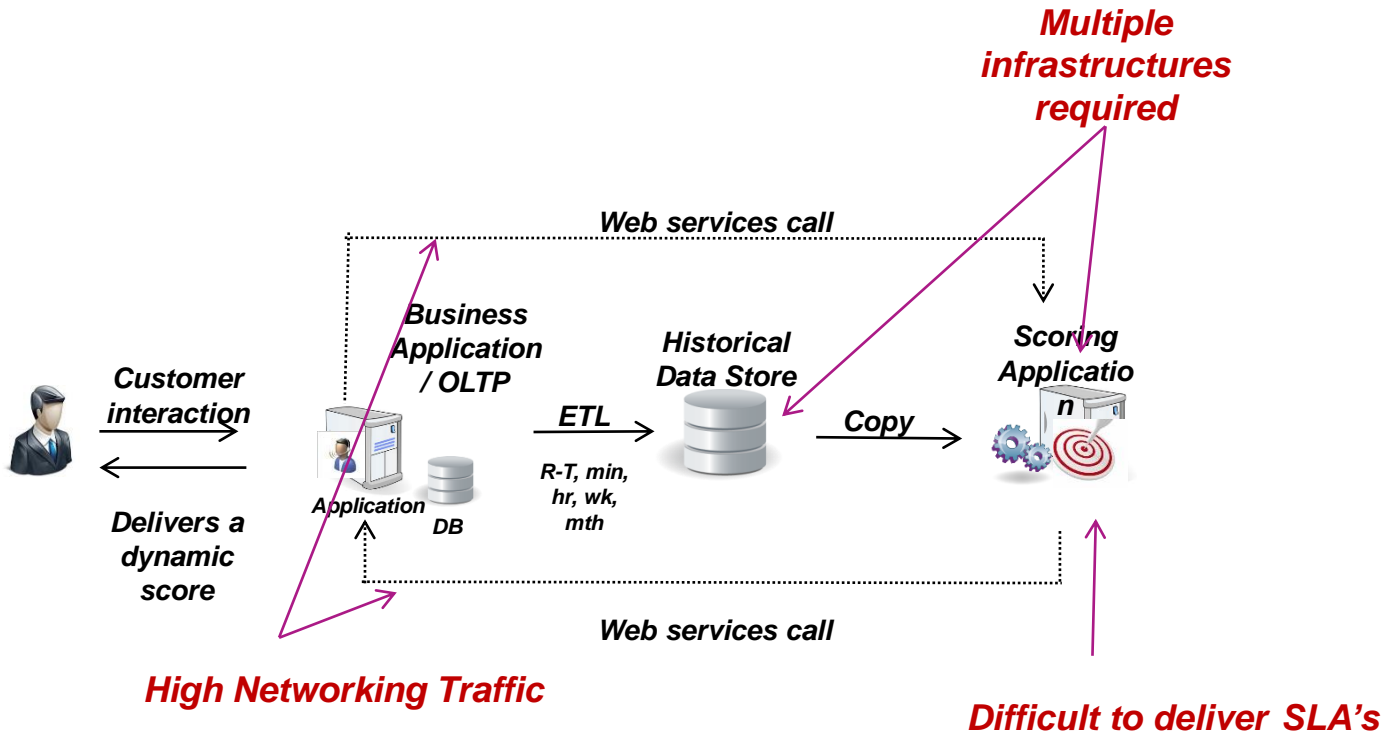
IBM SPSS Modeler Real-time Scoring with DB2 for z/OS

In transaction and in-database scoring on the same platform

Low latency with transaction workflow scoring

End to end solution on a single platform

Customer Interaction

Business System / OLTP

DB2 for z/OS Data Historical Store

Internal ETL

Transactions In

R-T, min, hr, wk, mth

DB2 for z/OS

Application w/latest data

Real-Time Score/ Decision Out

Copy

No external network Complexity, Latency

In transaction scoring

Rules Engine ODM

SPSS Modeler

Automated Model Refreshes

Meet or Exceed SLAs

On-Line Transaction and Analytics Processing (OLTAP)

For Linux on System z

Scoring Models

Scoring embedded in transaction processing workflow

Consolidated, Shared Resources

Preliminary Performance Findings: End-to-End Detection

3 Predictive SPSS Models

Orchestration

(95 inputs)
Demographic Anomaly   8%

Session Anomaly
(31 inputs)   10%

zHubAdapter

Fraud Propensity
(157 inputs)   9%

35%

16%

22%

FraudRule

doTransfer

Business Rule

Sample Transaction: CICS

z/OS Transaction Environment

Websphere CPU Cost			
zHubAdapter	0.00939		
doTransfer	0.00591	excludes CICS cost	
FraudRule	0.00412		
Fraud Propensity	0.00235		
Session Anomaly	0.00262		
Demographic Anomaly	0.00207		
Total	**0.02645**		

- Three advanced SPSS models executed during each transaction with many inputs

- Initial performance measures show favorable results (approx. 26.5 msec CPU time, end to end)

- More optimizations possible: z196 used, interface optimizations, etc.

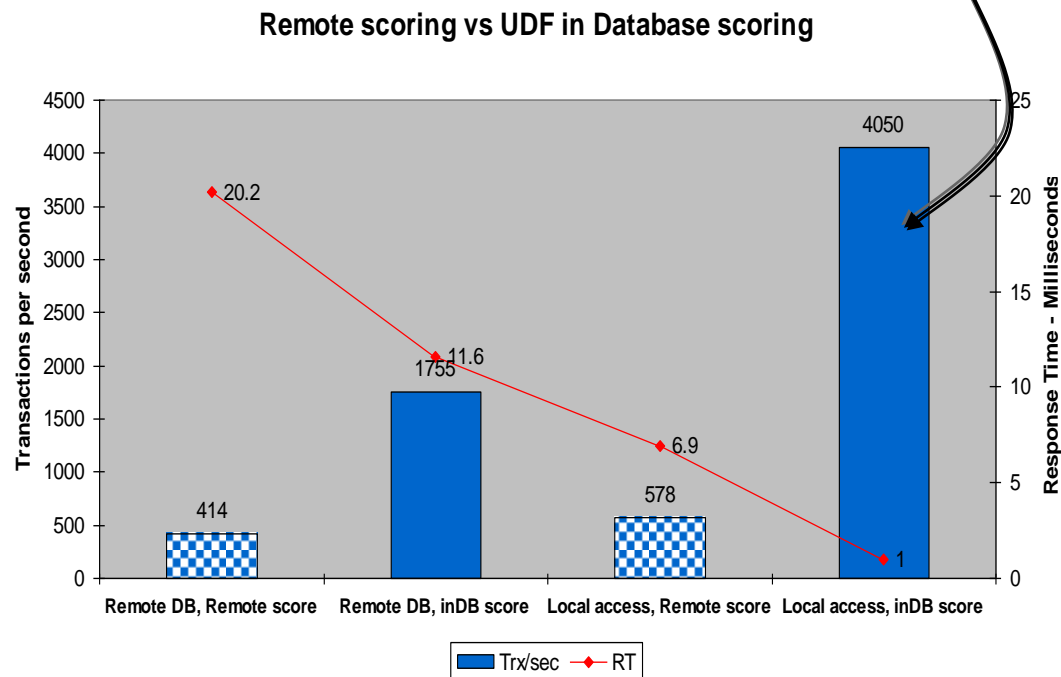# Compare Options: Scoring co-located with data on z vs. Remote Scoring

Co-locate analytics with data!

System z SPSS Scoring Performance
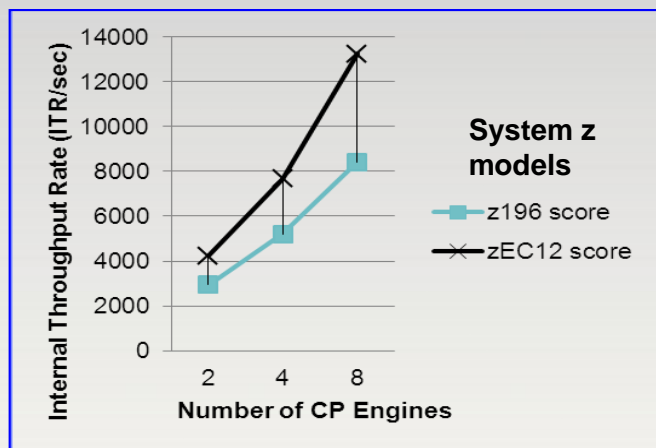
Execute thousands of model scores per second with very high scale

Unique: execute predictive models inside transactional database, with little data movement
- 7x performance improvement compared to moving data for analytics
- Achieve huge scales of execution without performance degradation
- Leverage historical and current transaction data to produce most accurate results

Data on z/OS: Orchestration and SPSS Score on z/OS

Remote scoring vs UDF in Database scoring

Column Terms	Definition
Remote DB	Application and DB2 z/OS are not collocated
Remote score	Scoring done by SPSS zLinux, data transferred from DB2 z/OS to SPSS via webservices
inDB score	Scoring done by DB2 z/OS with real-time scoring feature, no data transfer
Local access	Application and DB2 z/OS are collocated

Zementis for IBM z Systems (IBM Partner Application)
If you have other modeling applications, you can deploy the models on z Systems for Real Time Analytics!

- Zementis for IBM z Systems offers an integrated predictive analytics deployment and scoring capability for organizations managing data and transactions on z Systems

- Enables data scientists to retain full flexibility to develop predictive models in whichever commercial or open-source data mining tools they choose(e.g.IBM SPSS®, KNIME, Python, R and SAS), based on the PMML industry standard

- Helps rapidly deploy and easily integrate scalable, standards-based predictive analytics directly into critical business processes and execute predictive analytics within those processes.

- Delivers certified execution engines for applications compatible with the Java environment for z/OS (e.g.CICS®, WebSphere®)

- Scenarios such as fraud risk detection and customer insight analytics are supported in Zementis for z Systems

- Delivers a lightweight Java implementation for predictive analytics that executes on IBM z Systems zIIP specialty engines

Operational Decision Manager for z/OS: The Rules Engine

- Improve agility, productivity, and consistency by centralizing business rules and decision management

Without ODM

Application

Decision logic

With ODM

Application

Operational decisions
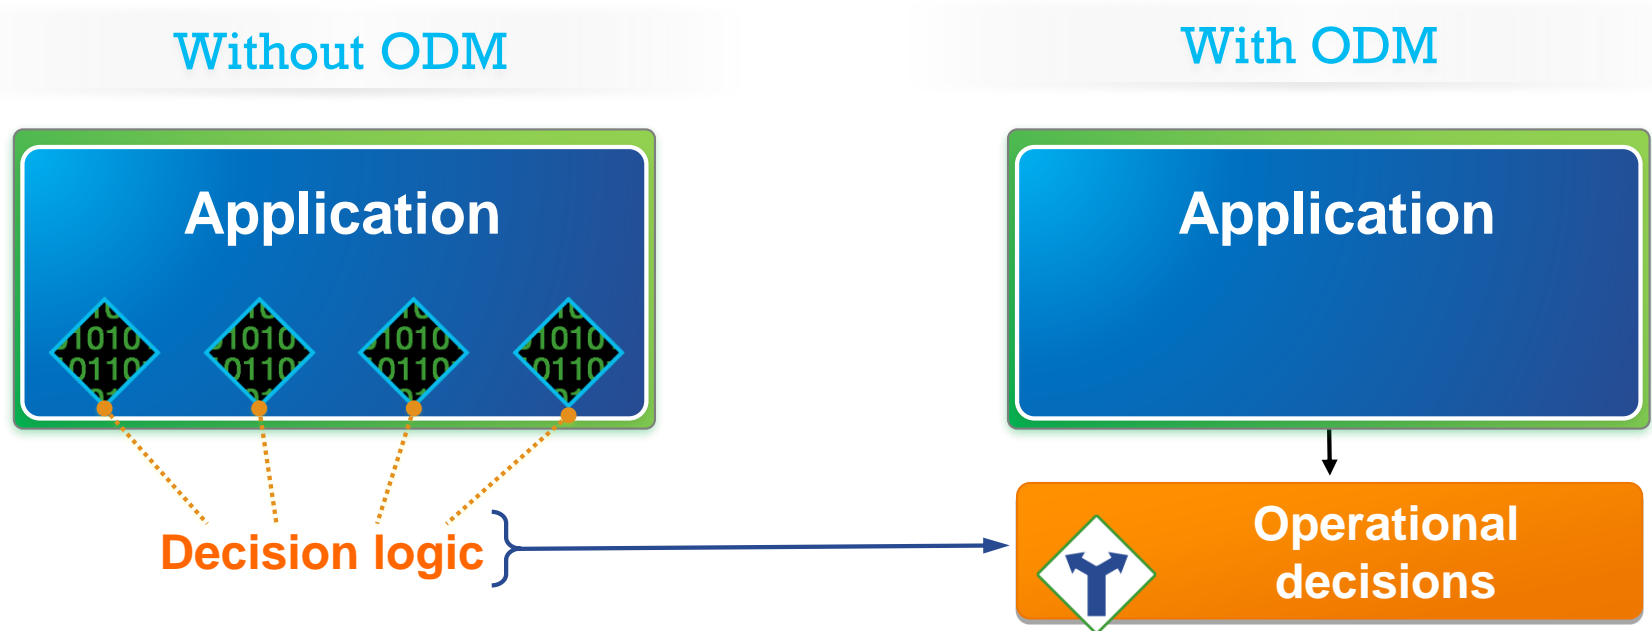
- Hard coded decisions: difficult to change
- Impedes rules re-use by other systems

- Externalized decisions: easy to change
- Centralized decisions enable reuse and consistency

The DB2 Analytics Accelerator
Speed up Predictive Model Building, Accelerate Complex Queries & Provide Vast Near Line Storage
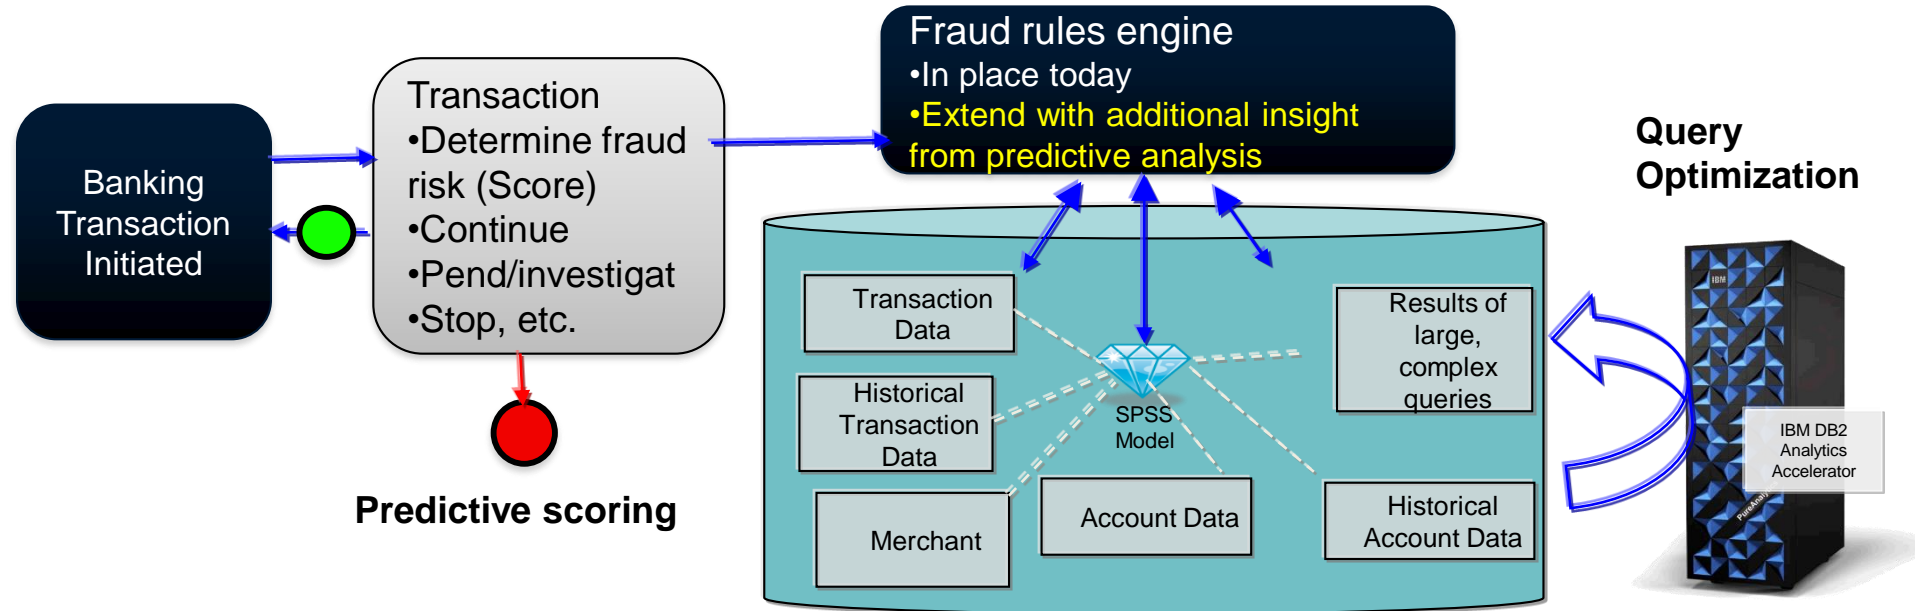
An <u>extension</u> of a DB2 for z/OS <u>system</u>, versus an "appliance"

- Transparent – queries are directed to a front-end DB2 system, as usual

- Self-managing – the optimizer of the front-end DB2 system determines where a query should be executed, and routes accordingly

- Highly integrated with front-end DB2 system, including messaging and monitoring

- Highly secure: the Accelerator inherits data access controls of front-end DB2 system

- Extreme acceleration of complex and/or data-intensive queries
  - 2000X faster (or *better*) in some cases
  - Response times can go from hours to seconds

Use Case 1:
Countering fraud, waste, abuse and financial crimes
(Banking example)

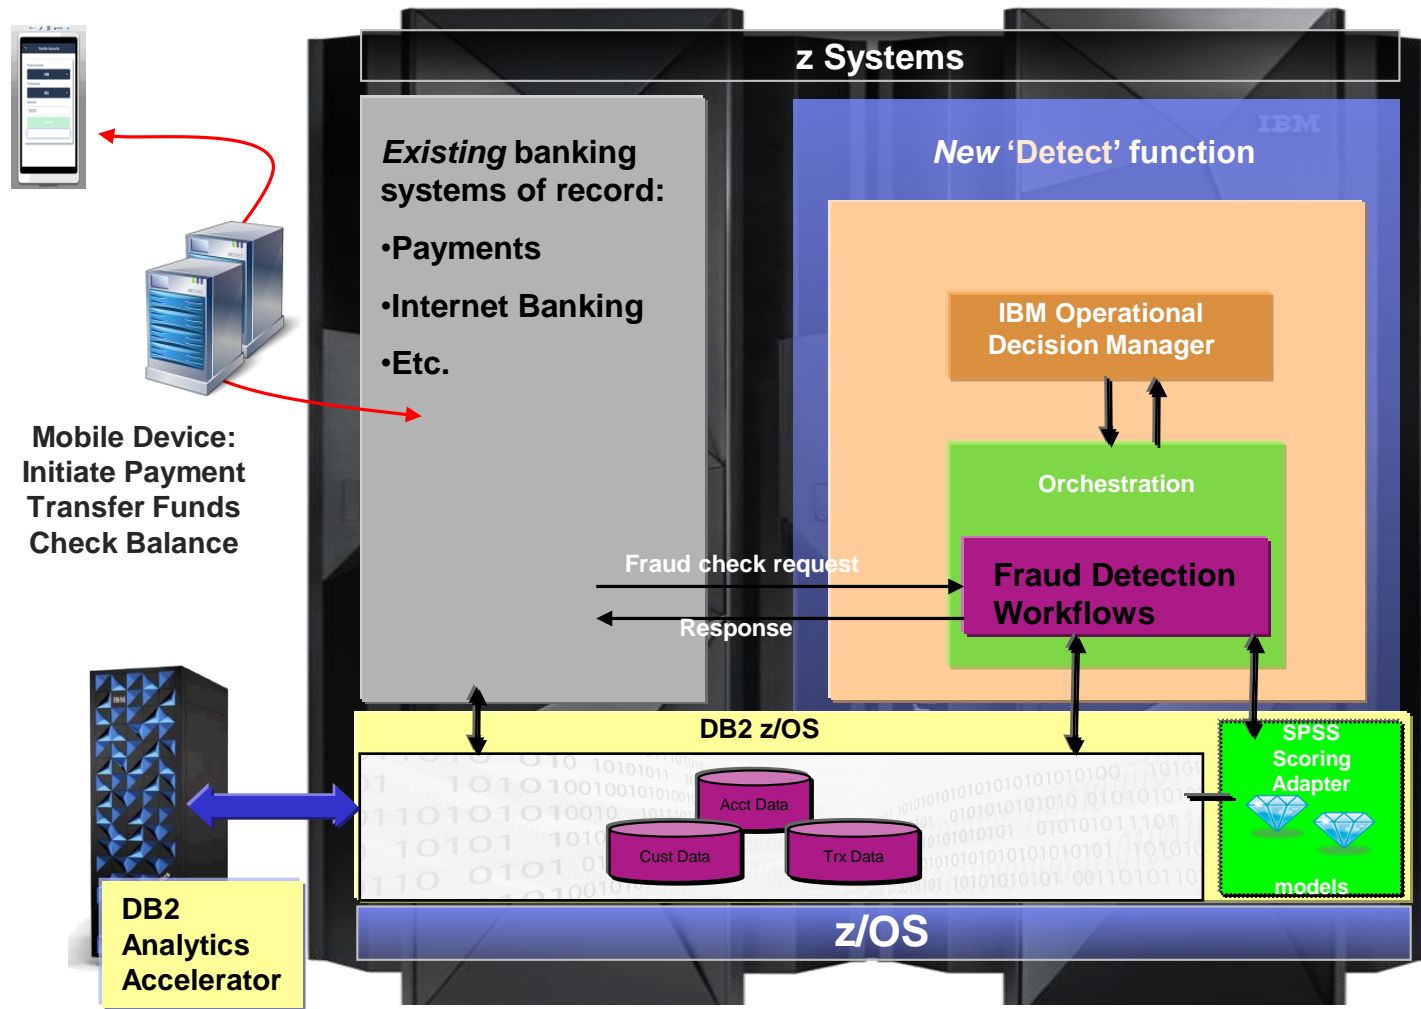# Business view - Bank: *Enhance* Existing Fraud Detection With Predictive Capability

Banking Transaction Initiated

Transaction
•Determine fraud risk (Score)
•Continue
•Pend/investigat
•Stop, etc.

Predictive scoring

Fraud rules engine
•In place today
•Extend with additional insight from predictive analysis

Transaction Data

Historical Transaction Data

Merchant

SPSS Model

Account Data

Results of large, complex queries

Historical Account Data

Query Optimization

IBM DB2 Analytics Accelerator

Business Goal: Reduce loss due to card fraud, reduce card deactivations, grow revenue associated with card purchases, reduce call center costs, improve service yielding preferred card usage
Approach: incorporate aggregate data from geographic location, merchant, issuer and card history into existing card authorization business flow to reduce fraud while preserving transactional SLAs

- Integrated **high performance query optimizations** enable client to aggregate data several times a day and use this complex data as part of real-time fraud detection process
- Enhance with predictive scoring integrated with fraud detection transaction for even more preventive capabilities

Architecture View - Bank: Enhance Existing Fraud Detection With Real Time Predictive Capability
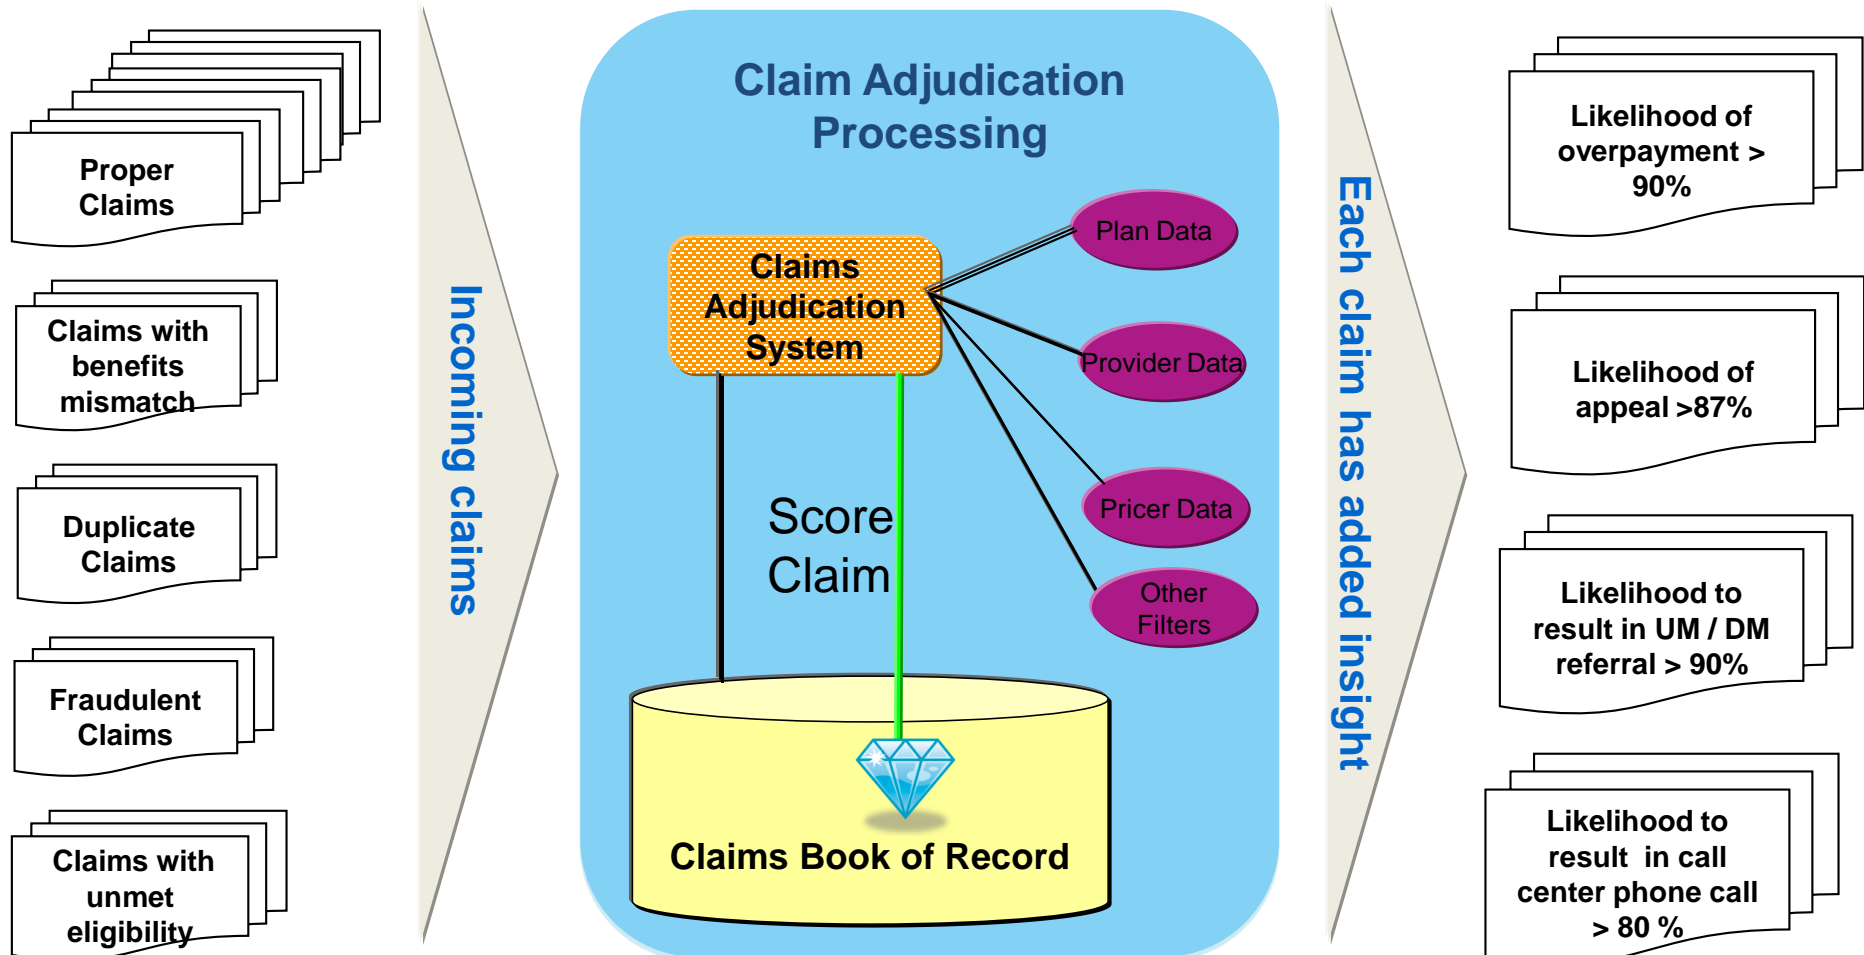
z Systems

Existing banking systems of record:

- Payments
- Internet Banking
- Etc.

New 'Detect' function

IBM Operational Decision Manager

Orchestration

Fraud Detection Workflows

Fraud check request

Response

SPSS Scoring Adapter

models

DB2 z/OS

Acct Data

Cust Data

Trx Data

z/OS

Mobile Device: Initiate Payment Transfer Funds Check Balance

DB2 Analytics Accelerator

Benefits

Fraud detection rates improved through integrated infrastructure that delivers new capabilities without sacrificing SLAs or compromising security

- Flexibility to invoke processes in real-time or scheduled, tight or loose coupled to invoking application

- Prioritize detect operations based on inputs (tran, LOB, amounts, etc)

- Leverage same infrastructure across multiple functions with both standardized processes & variations for specific LOBs as needed

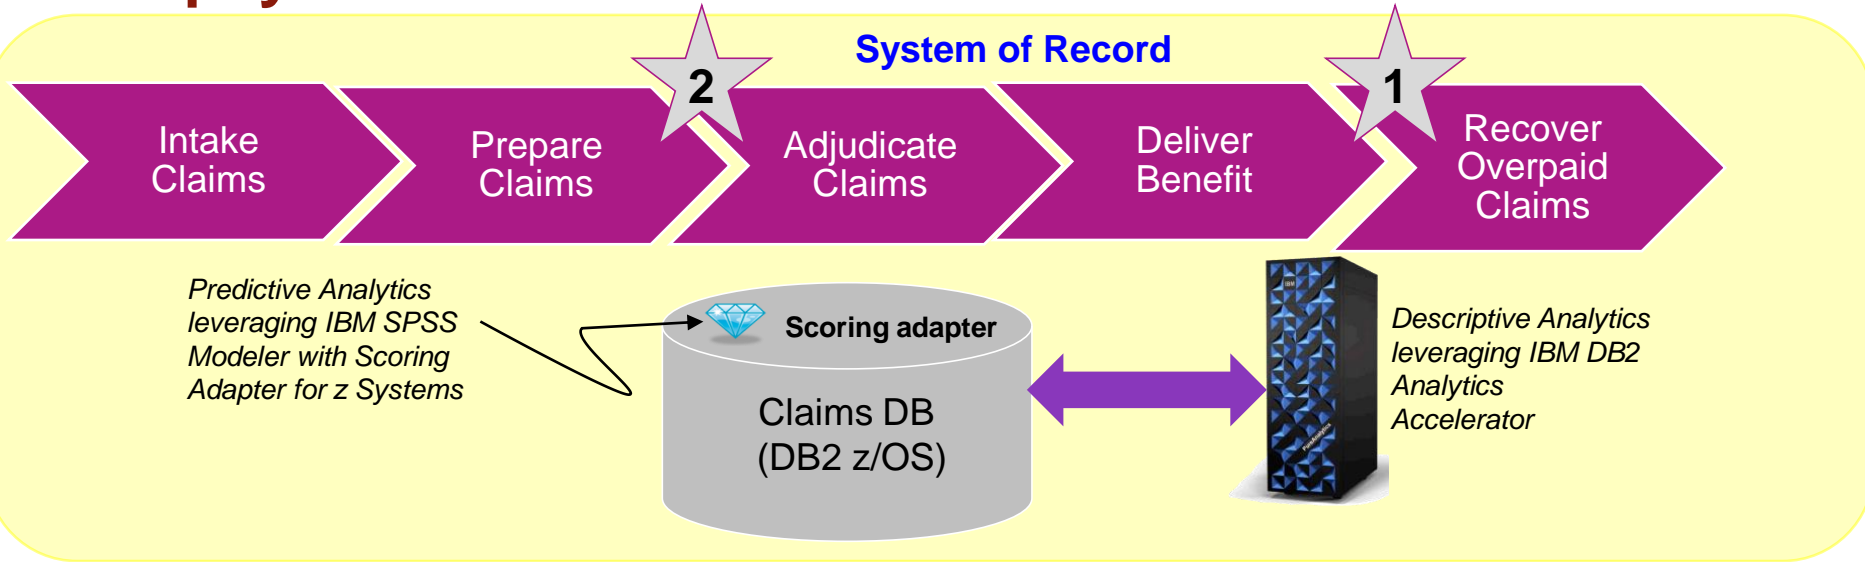- Can provide as a service across multiple clients, differentiate through context

Use Case 2:
In-process payment analytics
(Insurance claims example)

How can insurers leverage in-process scoring for best advantage?

Incoming claims

- Proper Claims
- Claims with benefits mismatch
- Duplicate Claims
- Fraudulent Claims
- Claims with unmet eligibility

Claim Adjudication Processing

Claims Adjudication System

- Plan Data
- Provider Data
- Pricer Data
- Other Filters

Score Claim

Claims Book of Record

Each claim has added insight

- Likelihood of overpayment > 90%
- Likelihood of appeal >87%
- Likelihood to result in UM / DM referral > 90%
- Likelihood to result in call center phone call > 80 %

Quickly, efficiently tag each claim with additional business insight; by extending in-process claims scoring across domains, you transform predictive analytics from a specialized function to one that best leverages transactional systems at a repeatable enterprise scale.

Business view - insurer: minimize loss from claims overpayments

System of Record

2

1

Intake Claims → Prepare Claims → Adjudicate Claims → Deliver Benefit → Recover Overpaid Claims

Predictive Analytics leveraging IBM SPSS Modeler with Scoring Adapter for z Systems

Scoring adapter

Claims DB (DB2 z/OS)

Descriptive Analytics leveraging IBM DB2 Analytics Accelerator

Challenge 1: Complex reports for overpaid claims not completing on time, result is monetary losses

Solution: Integrate optimized analytics of IDAA with overpayment reporting of transactions

Benefits:
- Up to 2000x improvement in speed of overpayment reports
- LOB users enabled to respond with more agility to overpayment trends
- Informed decisions at the right time

Challenge 2: Stop improper payments *prior* to payment, avoid pay & chase, meet SLAs

Solution: Integrate predictive analytics into claims adjudication for analytics in place

Benefits:
- Very efficient scale for analytics
- Scale requirements only achievable with analytics as part of transaction flow
- Expected results of efficient in-transaction analytics can be multi-million dollars per year

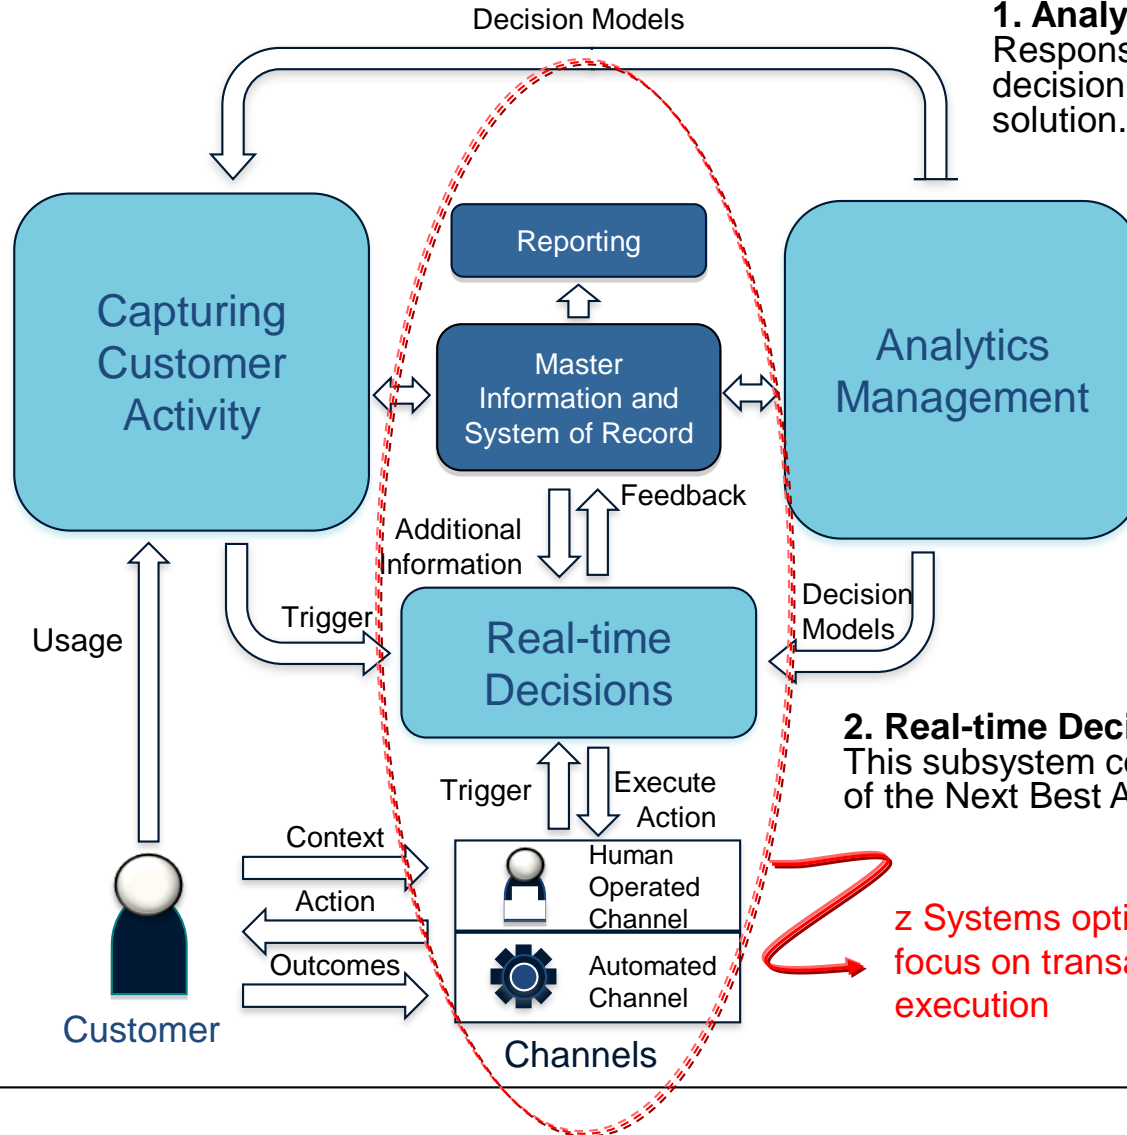# Insurance: Minimize Losses From Overpayments Architecture view

Incorporate model execution into z/OS claims transaction systems

MQ: Tran info

Get claims transaction information

Determine likelihood of overpayment

Use model result, determine best action

Recommendation Examples:
•Continue pay
•Pay but highlight for investigation
•Execute further models

MQ: Tran info

Invoke data preparations required for input to models

Invoke score: result is Propensity of Overpayment and confidence

- Efficient, scalable scoring on each and every claim

- No data movement to achieve analytics

- Leverage both historical and current transaction data

Current Transaction Data

Slow Changing Data

SPSS Model

Historical Data

Summarized Result Tables

DB2 zOS v10 or above

Use Case 3:
Predictive customer intelligence / executing the next best action

Next Best Action: High-level logical view

1. Analytics Management:
Responsible for developing the decision models for the solution.

3. Capturing Customer Activity:
This subsystem gathers information about the customer's activity, including the outcomes of the Next Best Actions offered by this solution.

Decision Models

Reporting

Capturing Customer Activity

Master Information and System of Record

Analytics Management

Feedback

Additional Information

Usage

Trigger

Real-time Decisions

Decision Models

2. Real-time Decisions:
This subsystem contains the execution of the Next Best Action decision loop.

Trigger

Execute Action

Context

Action

Human Operated Channel

Outcomes

Automated Channel

z Systems optimized NBA functions focus on transactional, real-time execution

Customer

Channels

Architecture view: Real-time analytics for next best action

z Systems

Existing z/OS systems of record:
- Transactions
- Payments
- Orders
- History

New z/OS-optimized NBA function

IBM Operational Decision Manager

Orchestration

NBA Workflows

Trigger / Execute Action

Sensors producing readings

Internet

Social Media Services

Customer Conversations

Capturing Customer Activity

Process
Process
Process
Process
APPLICATIONS

Files

Shared Staging Area

CEP

Usage

Context
Action
Outcomes

Customer

Human Operated Channel
Automated Channel

Channels

DB2 z/OS

Acct Data
Cust Data
Trx Data

SPSS Scoring Adapter

models

z/OS

DB2 Analytics Accelerator

Business view - Bank: Optimized Next Best Action for Customer Initiated Loan

Short tenure small business LOC requesting car loan

Customer call: need a new business car loan

Advise: eligible for new car loan up to $35,000 with $1,000 down on vehicle

Eligible to leverage current LOC for car loan with improved overall rate and LOC credit increase. If LOC used for loan, then down payment not required

Offer accepted by customer

Offer decision considerations

Capturing Customer Activity

Master Information & Systems of Record Transactions

Real-Time Analytics & Decisions

History

Transactions

Customer

Model

Loan Payments
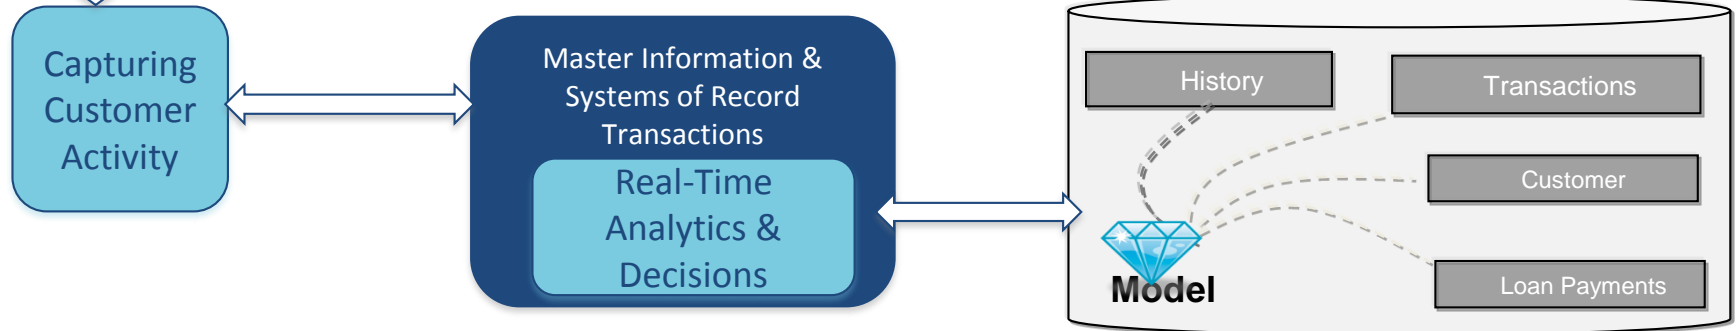
- Banking clients want to integrate information across all product lines in order to make **real-time, targeted decisions**

- Why real-time? Bank may risk losing customer business or loyalty to other products

- Need to incorporate high value, predictive advanced analytics as part of transactional systems

- Why do we need all the data to score? Reduce bank's risk in approving loan

Real-time predictive customer intelligence leads to smart business decisions at the point of impact

| 1. Build long-term customer relationships | 2. Drive one decision one interaction at a time | 3. Maximize customer life-time value |

Telco Providers

- When a high-value, long-tenure **customer calls with a bill query** within six months of the end of his / her contract, **offer** him / her a **complimentary handset upgrade** if he / she renews in advance.
- Customer **repeatedly called tech support** with mobile Internet issues comes to the website to check process to transfer his / her number. Initiate a chat, check problems are now resolved, **refund** two months data charges by way of apology, and **offer discount** on a new handset.

Banking Providers

- **Customer calls** for password reminder on "self-trade" account. Pattern of usage shows limited activity and few returns; **offer upgrade** to managed investment offering.
- Customer's transactions show **multiple trips to Asia; offer** international emergency cover, and provide information about commission-free ATM withdrawals from our regional partner.

Insurance Providers

- When confirming to a customer that his / her claim will be settled, **direct him / her to a repair shop** which, although it may take 1 – 2 days longer than other options, has higher quality ratings.
- When a **customer calls** in with a coverage query, let them know **they can save money** by combining their two policies into a single multi-car one.

Real-Time Scoring Use Cases:

Use Case	Industry
Credit card fraud detection/prevention	Banking
Risk and claims fraud detection/prevention	Insurance
Risk, crime, fraud detection/prevention	Government
Cross sell, up-sell from market basket analysis	Retail
Cross sell, up-sell from market basket analysis, scheduling, inventory, irregular operations	Travel
Optimize logistics from telemetry, fuel consumption, schedule and weather patterns	Transportation
Churn management, IVR/Call Center Processing	Telco
Meter management, Asset management	Utilities

Mobile Analytics & z Systems in Action

IBM CFM

Mobile Device
Initiate Funds Transfer Request

Suspicious Transaction Management

- **Link Analysis**
- **Alert / Case Management**
- **Additional Investigation**

zLinux
Worklight Server

Suspicious Trans

z/OS Transaction Environment

| Prepare Data for Predictive Models | Invoke 3 Predictive Models | Prepare Data for Business Rule | Invoke Business Rule | Determine Fraud Risk Score | Risk of fraud high? | Queue Suspicious Transactions | Return Status and Response |

Y → Queue Suspicious Transactions

N → Sample CICS Banking Transactions

SQL calls

ODM Business Rules Server

Non-Suspicious Transaction, do transfer

DB2 z/OS Data

Trx Data — **Models** — History — History — Trx Data — Acct Data — Cust Data

Analytics Accelerator regularly preparing data for predictive models

Complete transaction in 26ms

Results

- **Showed advanced analytic fraud detection executing at bank transaction speeds**
- **Three predictive models with large numbers of inputs (31, 95, 157) executed in real-time**
- **Low IT consumption – does not impact BAU workloads or MIPS**
- **Suspicious transactions queued for additional investigation and processing**

51

Real Time Analytics on System z with DB2 – Value Proposition

- Reduced latency for true Real Time Analytics; Limited data movement and copies

- Reduced ETL: Some clients have lowered ETL storage costs by up to 75%

- Benefit from high QoS: Availability, Integrity, Security, Scalability etc.

- Scoring workload has minimum impact on servers, networks and consequently on SLAs (IBM Tests)

- Much easier to incorporate scoring into other applications and work processes

- Scoring on most current data (last committed) and scoring in a single SQL statement; Score 100% of transactions (no need for sampling)

- Savings due to reduced storage and network components

- No additional data integrity and governance issues; Single view of enterprise data! 'Single version of the truth'

- No additional security exposures. Exploit System z's ultra high security;

- Refreshing the models is easier and immediate

Questions?

Need an on site briefing for your team?

Topics:

Analytics on z
Real Time Analytics on z
Spark on z/OS

Please contact:

Rajeev V Kamath in Raleigh at rvkamath@us.ibm.com

▪ **Benefits**

- Performance: tight, secure, audited microflows invoking orchestration, rules, scoring

- Flexibility to invoke processes real-time or scheduled, tight or loose coupled to invoking application

- Prioritize Fraud Detection operations based on inputs (tran, LOB, amounts, etc)

- Leverage same infrastructure to detect   across multiple areas

- Re-use *same* core reference architecture for AML capabilities