

Business-critical insight for proven competitive advantage

Gary Crupi, IBM Distinguished Engineer, IBM Analytics Client Advocate
June 2015



Goals for Today

- **Understand...**

- The Spectrum of Traditional and Emerging Analytical Capabilities
- The z Systems Advantage and Point of View

- **Act...**

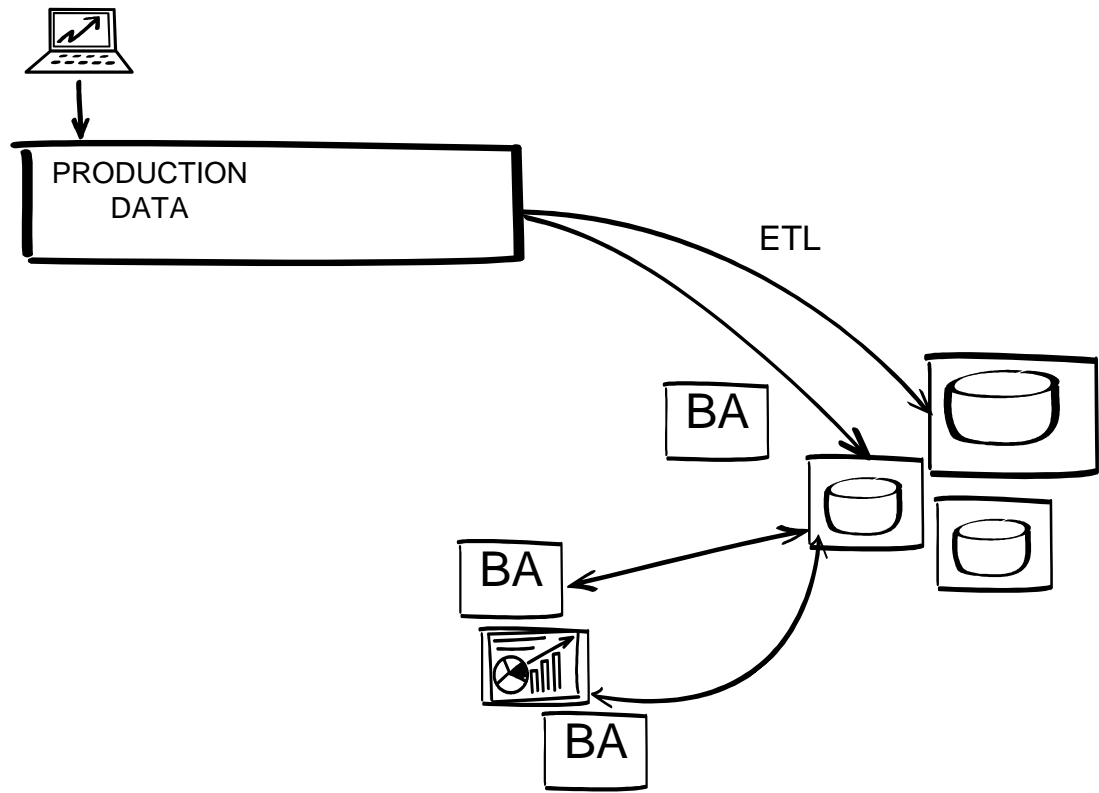
- Deepen your knowledge of the DB2 Analytics Accelerator
- Determine if you can leverage the z Systems Analytics Advantage
- Engage with IBM to determine the best path for adopting the approach

Spectrum of Traditional and Emerging Analytical Capabilities

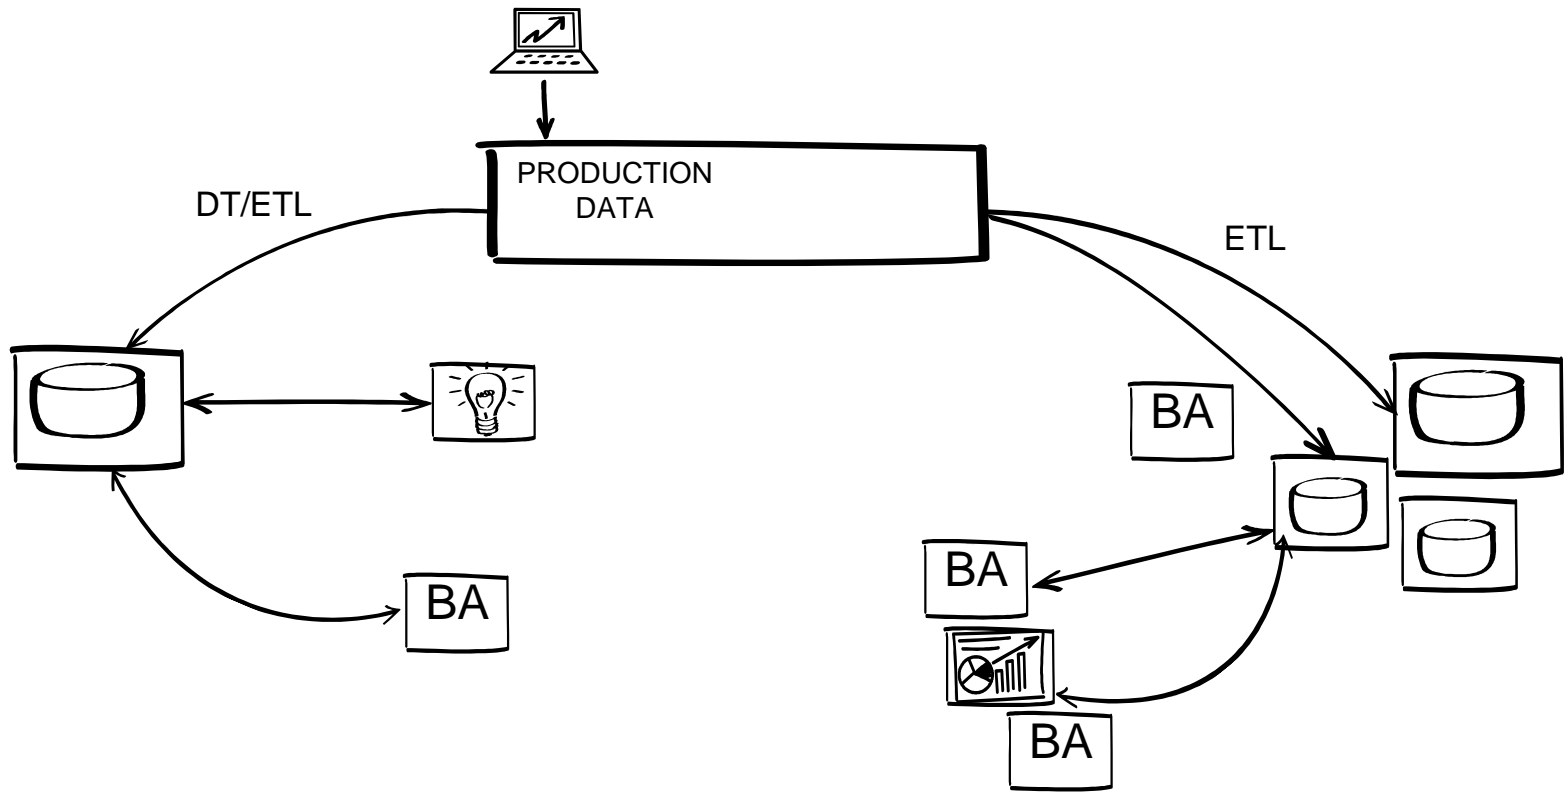


58% → Analytics → 220%


58% → Analytics → 220%



58% → Analytics → 220%

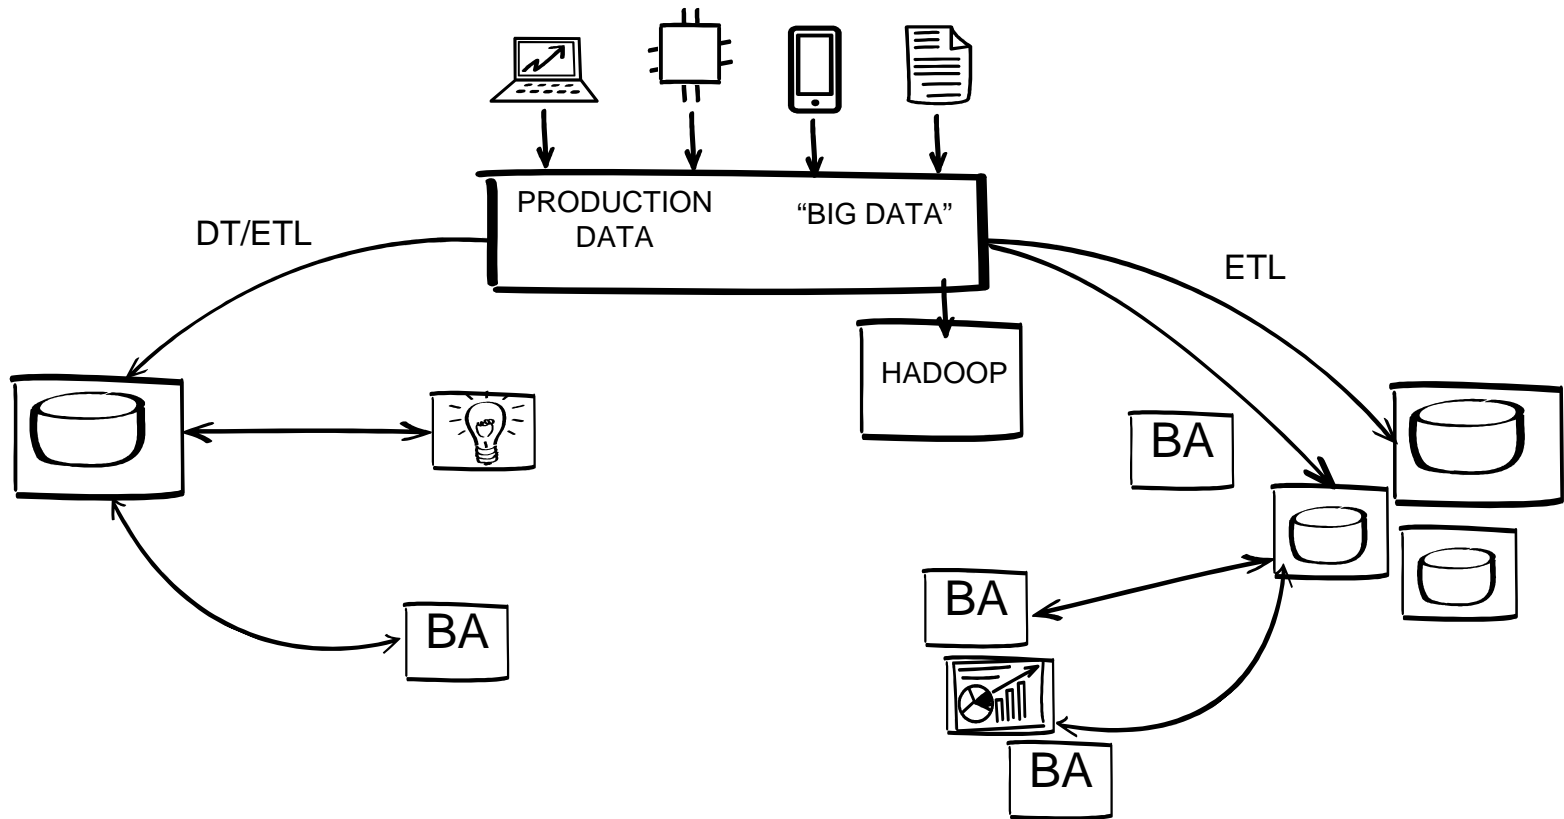


58% → Analytics → 220%



COMPLEXITY ↑

58% → Analytics → 220%



COMPLEXITY ↑

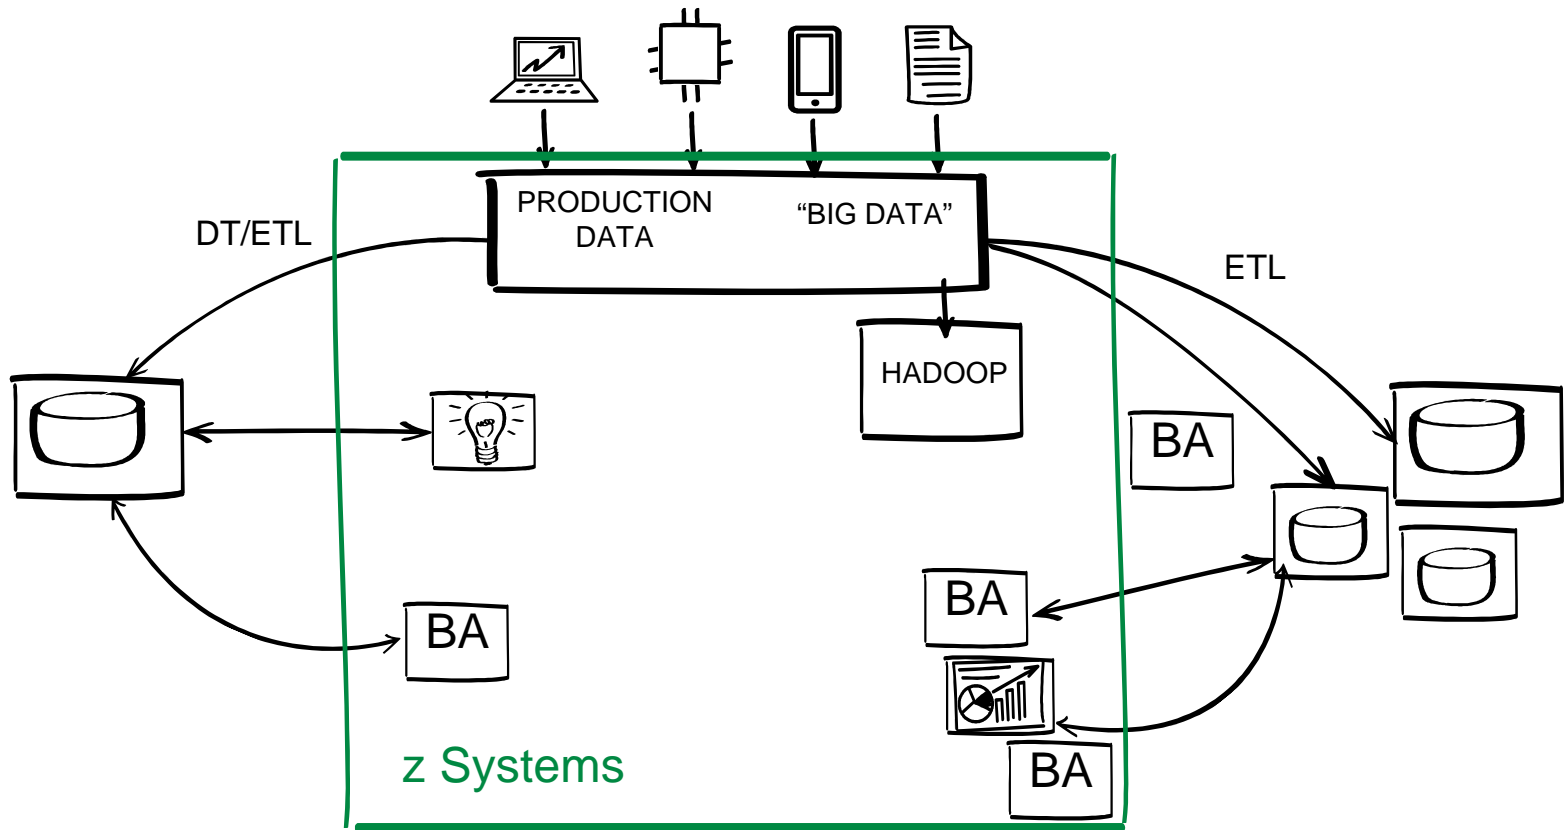
TIME

PERFORMANCE

RELIABILITY

SECURITY

58% → Analytics → 220%



COMPLEXITY ↑

TIME

PERFORMANCE

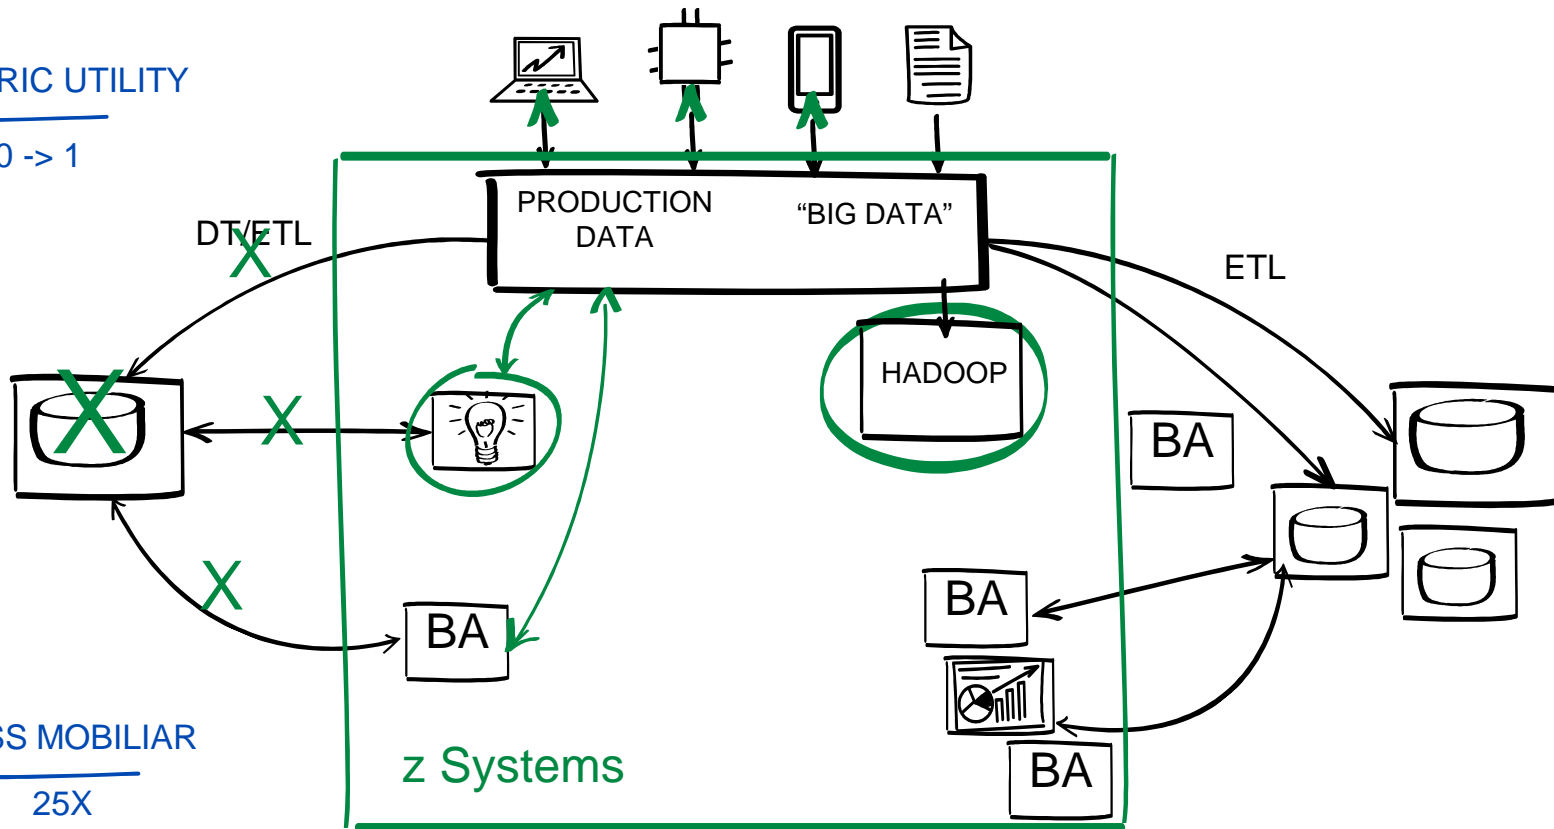
RELIABILITY

SECURITY

58% → Analytics → 220%

ELECTRIC UTILITY

10 → 1



SWISS MOBILIAR

25X
100X

COMPLEXITY X

TIME

PERFORMANCE

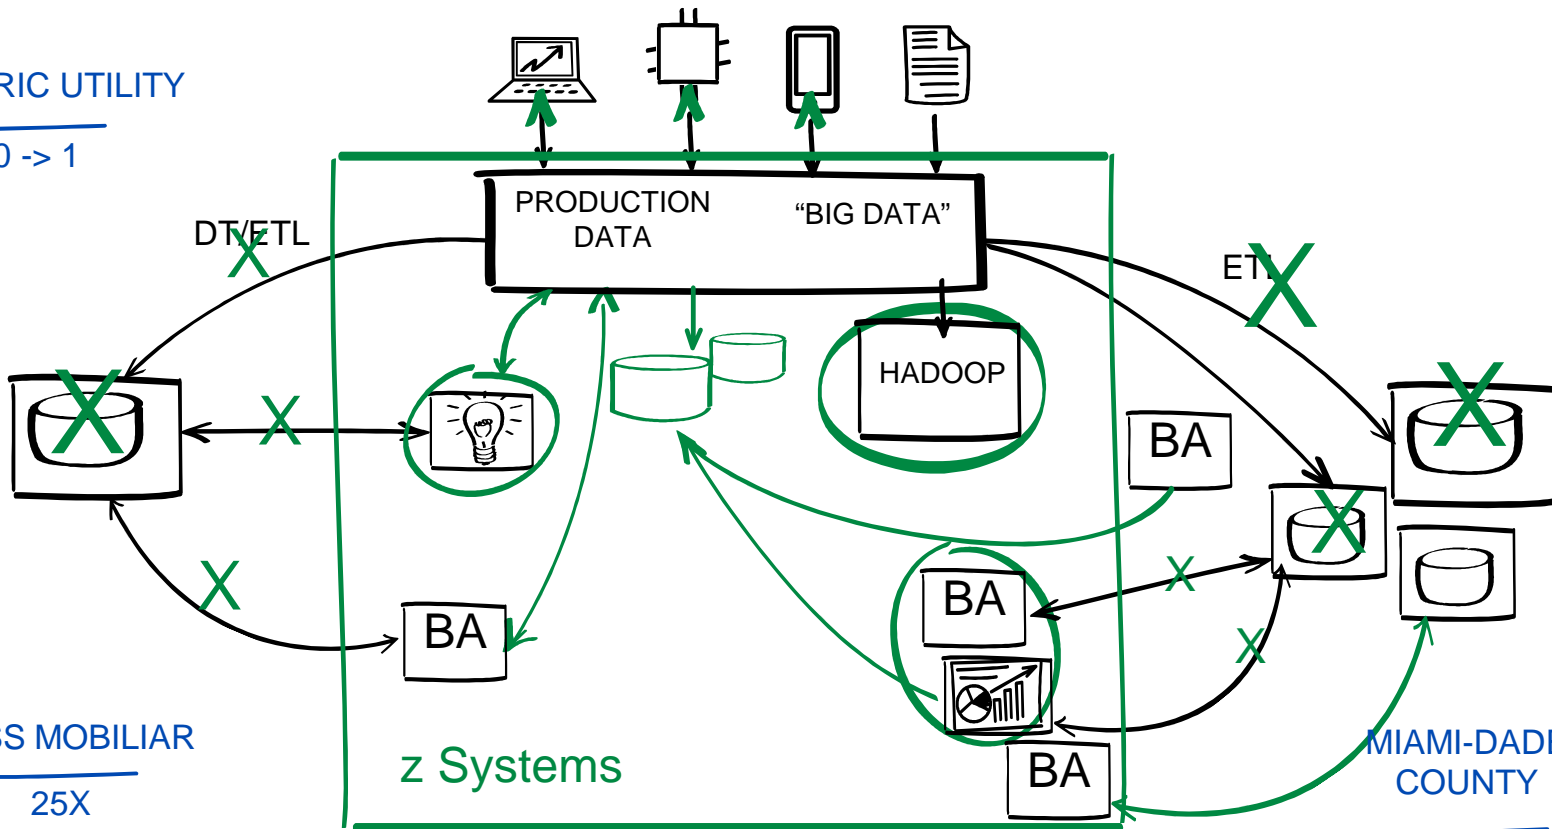
RELIABILITY

SECURITY

58% → Analytics → 220%

ELECTRIC UTILITY

10 → 1



SWISS MOBILIAR

25X
100X

MIAMI-DADE
COUNTY

11 DAYS

COMPLEXITY X

TIME

PERFORMANCE

RELIABILITY

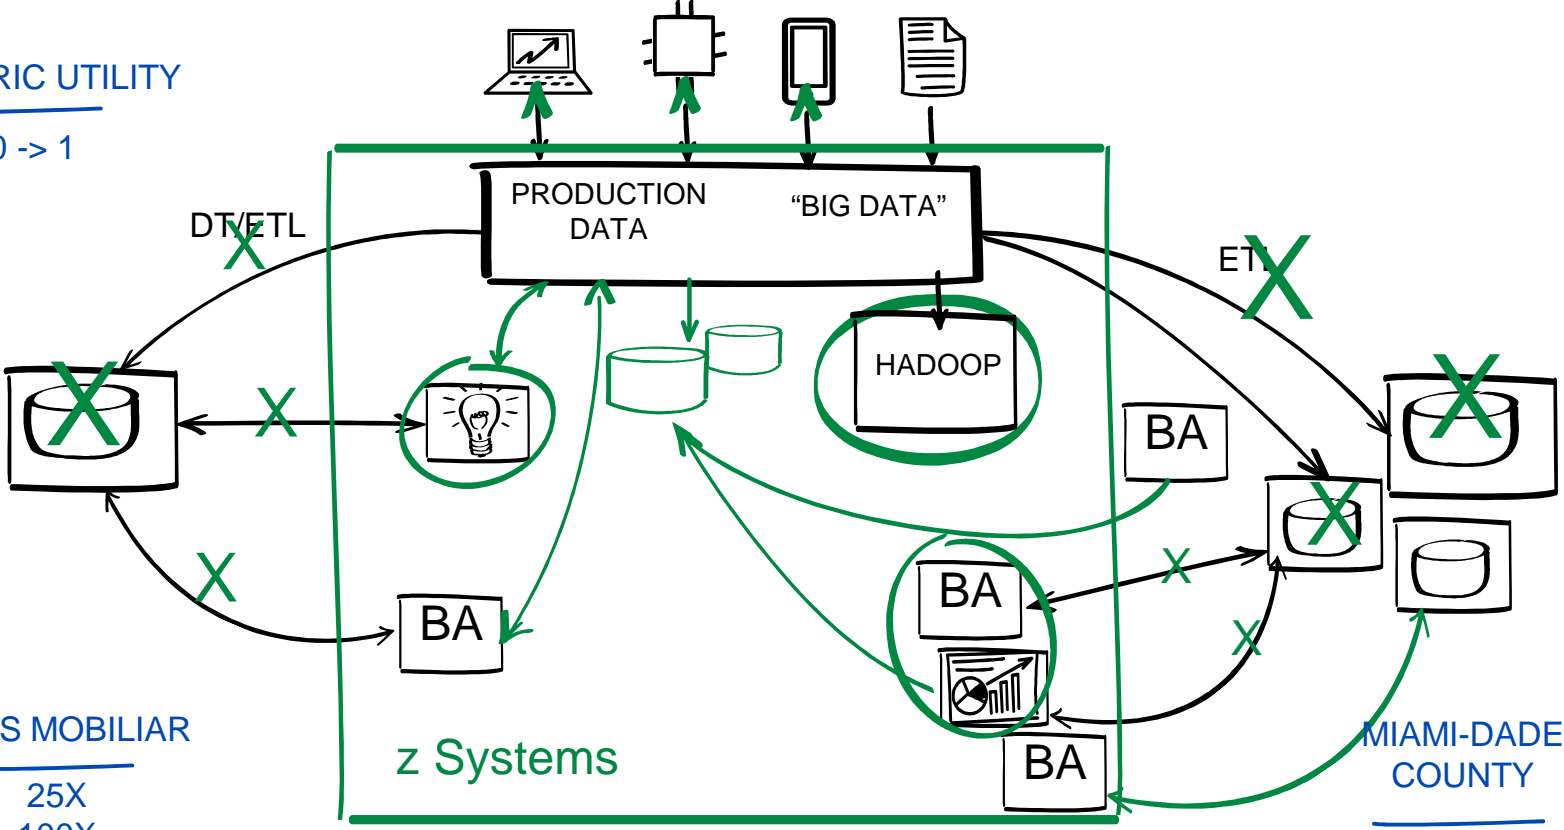
SECURITY

z Systems

58% → Analytics → 220%

ELECTRIC UTILITY

10 → 1



SWISS MOBILIAR

25X
100X

MIAMI-DADE COUNTY

11 DAYS

COMPLEXITY **X**

TIME ↓
24 → 3

PERFORMANCE ↑
2000X
12K TPS

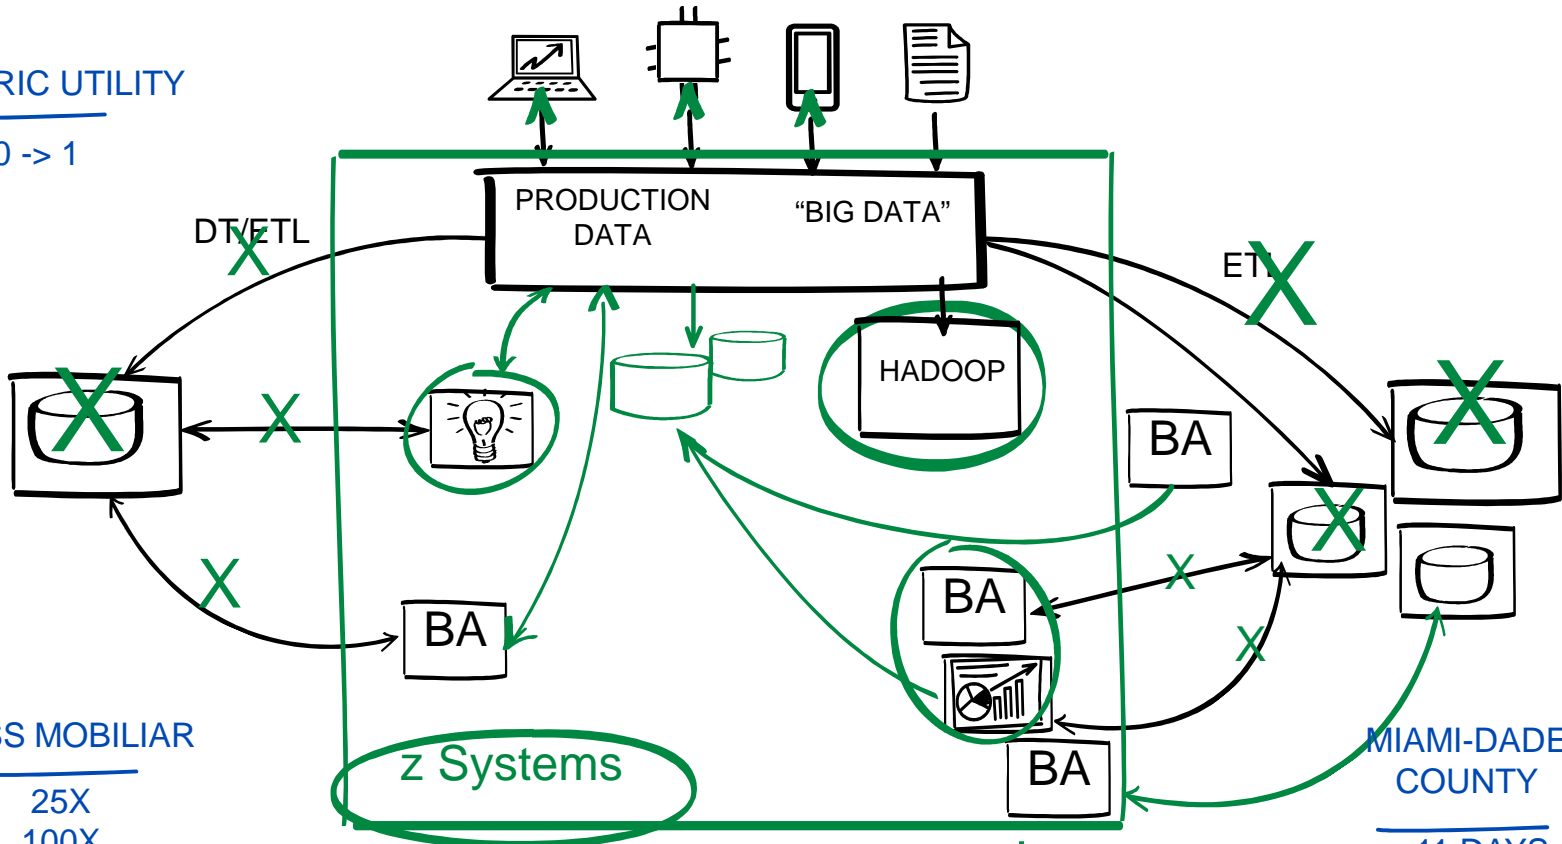
RELIABILITY ↑
"THE SOURCE"
99.999%

SECURITY

58% → **Analytics** → 220%

ELECTRIC UTILITY

10 → 1



SWISS MOBILIAR

25X
100X

COMPLEXITY X ↓ \$

TIME ↓
24 → 3

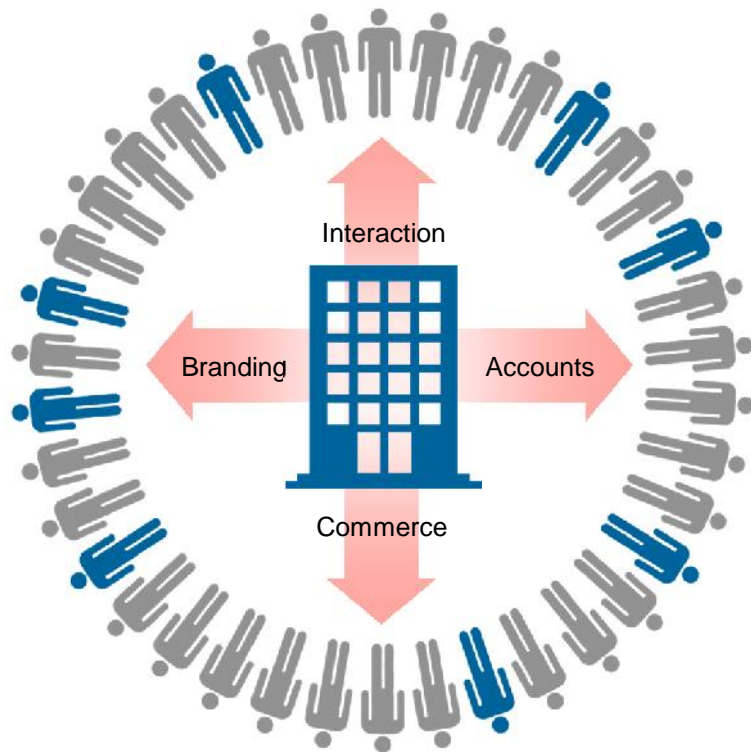
PERFORMANCE ↑
2000X
12K TPS

RELIABILITY ↑
"THE SOURCE"
99.999%

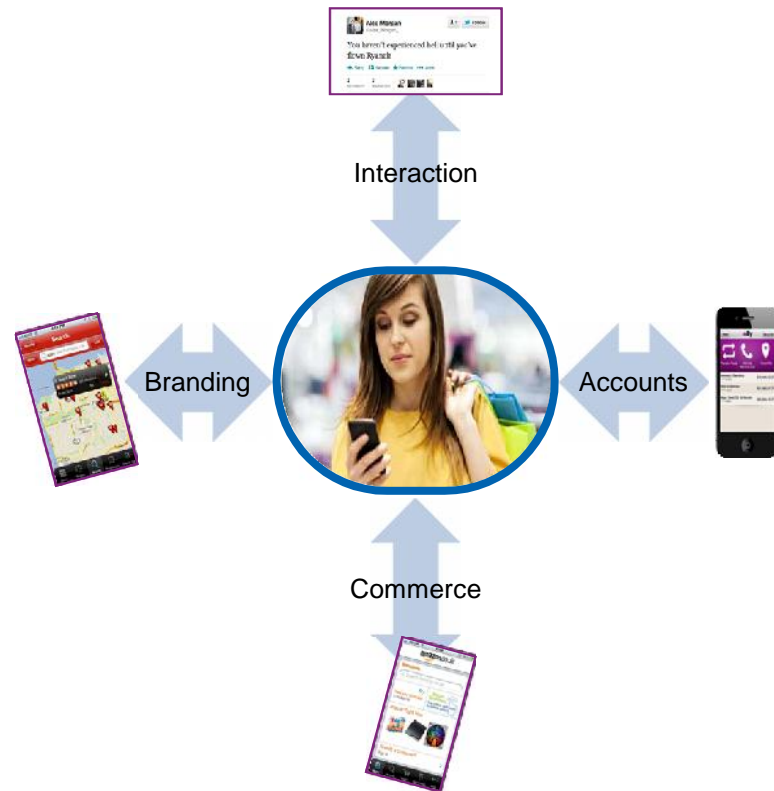
SECURITY


The business-customer relationship has changed

Then: *“I have an offer – let me find a customer I can sell to”*



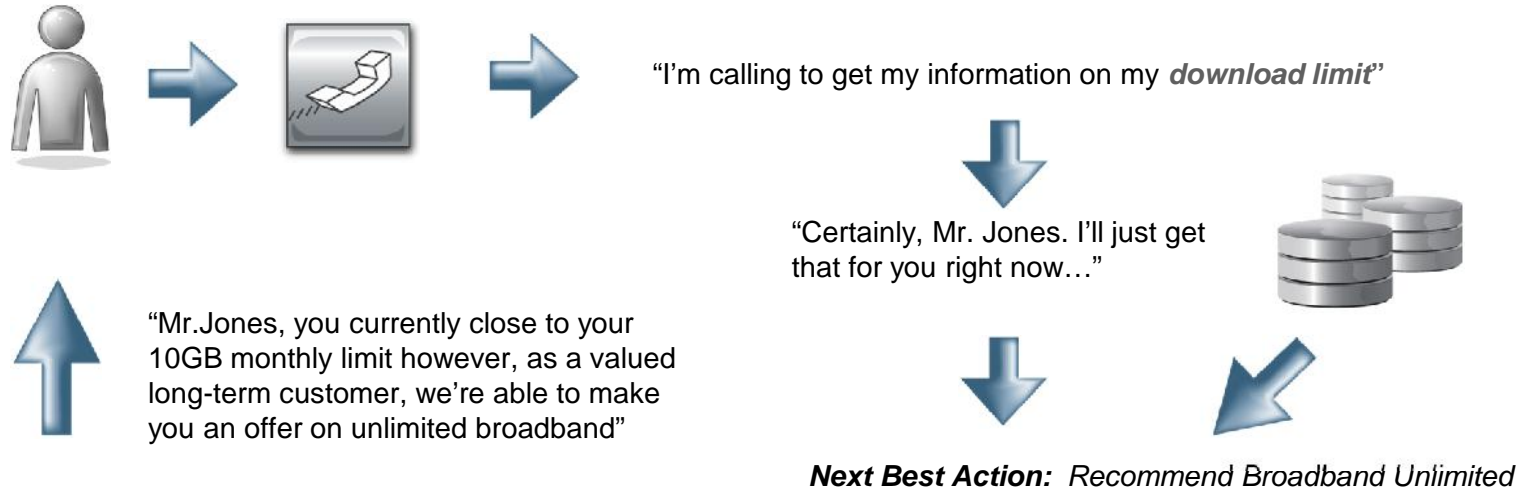
Now: *“I have a customer – what do they need most?”*



Customer experience is the competitive advantage for top-line growth

Real-time Analytics:

Maximize Customer Wallet Share & Improve Customer Experience



Building the business case

- Upsell =
 - \$10 increase/month = \$120/yr
- Call volume = 1Million calls/yr = 2777 calls/day
- Current success rates = 25% = \$30 Million
- Improve it to 35% = \$12 M increase in revenue

Campaign	Market Segment	Predicted Satisfaction	Eligible	Likelihood	Predicted ARPU	Predicted Margin
Retention	1	64%	NO	12%	10	2.3
Broadband Unlimited	2	68%	YES	68%	13	5.2
Anywhere Smartphone	2	68%	YES	72%	23	4.3

Real-time Analytics:

Banking: Reduce Fraud & Improve Customer Experience



Business Goal:

- Increase insight from merchant, geographic, store data
- Incorporate customer sentiment from social media & email
- Reduce false positives of card deactivation
- Reduce customer interaction costs

Business value:

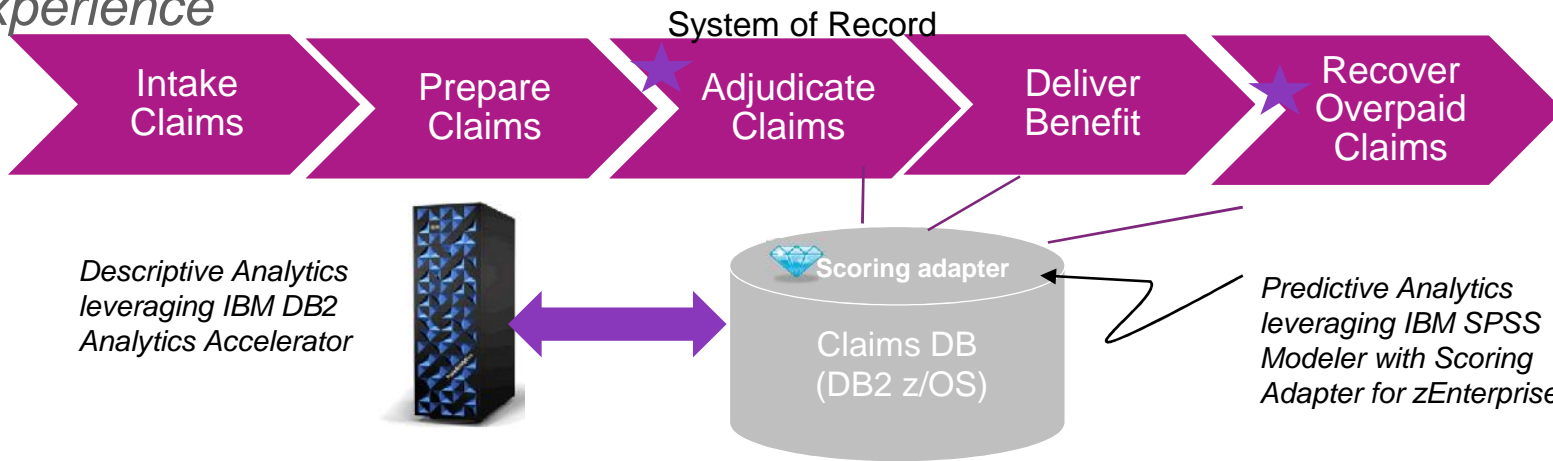
- Reduce loss due to fraud
- Improved customer satisfaction with personalized experience
- Drive incremental revenue with up-sell/cross sell
- Reduced call center costs
 - Less calls
 - Faster response time

Solution

- DB2 Analytics Accelerator to aggregate records more often than previously possible to learn from & apply new data
- SPSS Modeler & DB2 Scoring to analyze new transactions in real-time to proactively and accurately predict fraud
- Cognos BI to provide the front line with the analysis to maximize customer wallet share
- BigInsights to integrate, understand & proactively manage customer sentiment

Real-time Analytics:

Insurance: Minimize Loss from Claims Overpayments & Improve Customer Experience



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

Solution: Integrate optimized analytics of DB2 Analytics Accelerator 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

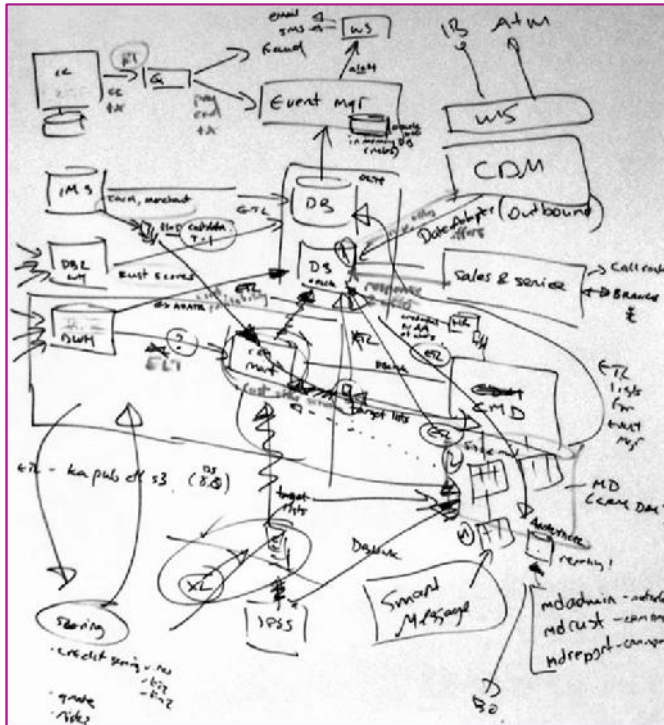
Benefits:

- 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

The z Systems Advantage and Point of View




Challenges with traditional analytics processing



- **Significant complexity**
Data is move from operational databases to separated data warehouses/data marts to support analytics
- **Analytics latency**
Transactional data is not readily or easily available for analytics when created
- **Lack of synchronization**
Data is not easily aggregated and users are not assured they have access to “fresh” data
- **Data duplication**
Multiple copies of the same data is proliferated throughout the organization
- **Excessive costs**
An IT infrastructure that was not designed nor can support real-time analytics

Transactions and analytics processed together



- Purchase made
- Resources consumed
- Bill paid
- Claim submitted
- Information updated
- Call center contacted


- What happened?
- How many, how often, where?
- What actions are needed?
- What will happen if?
- What will produce the best outcome?

Analytics as part of the flow of business; insights on every transaction

Integrated analytics and transaction processing

The greatest value is achieved when analytics are run where the transactions and data originate

Insights on every transaction



Analytics as part of the flow of business



More transaction throughput

Faster data access

Increased analytics performance

Avoid (ETL)

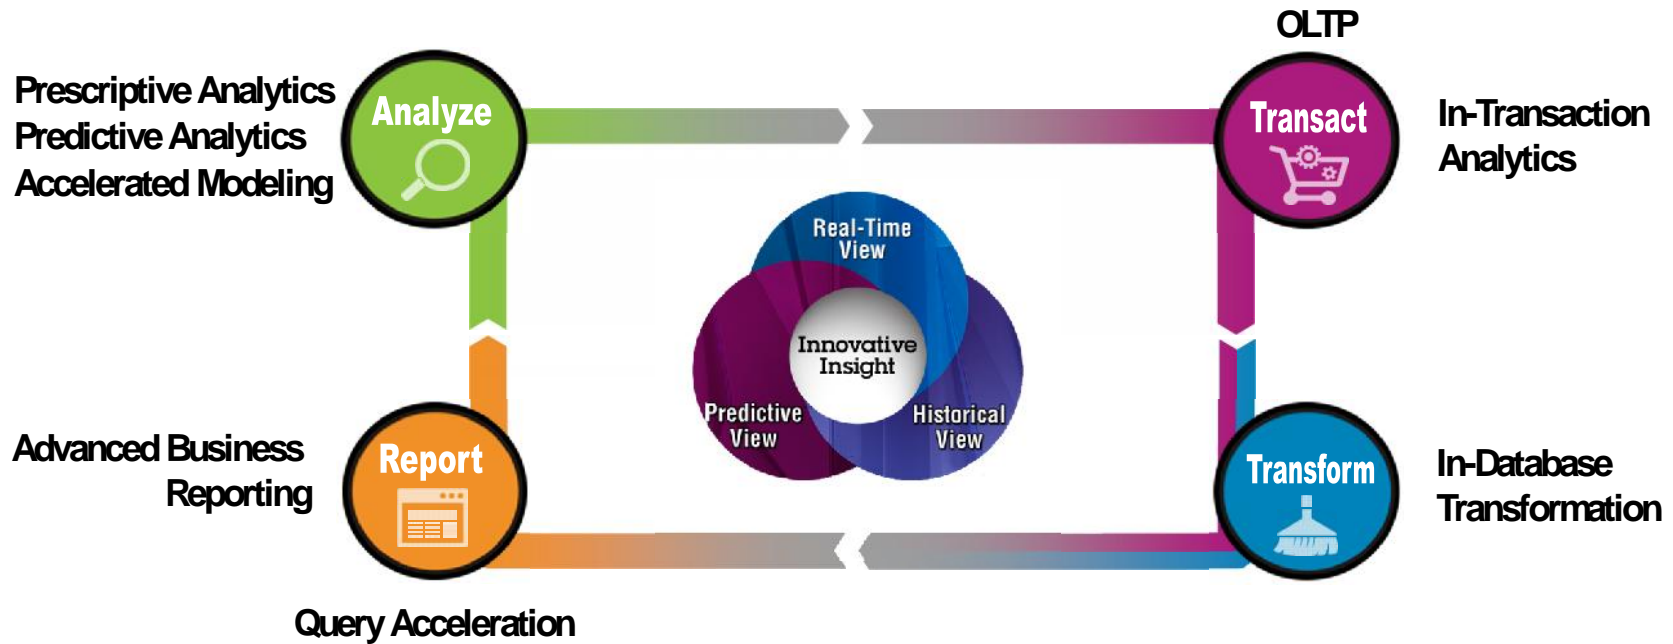
- Eliminate complexity
- Eliminate cost
- Eliminate redundancy

Deliver real-time insights at the point of impact

- Improve customer experience
- Increase revenue opportunities
- Improve response time

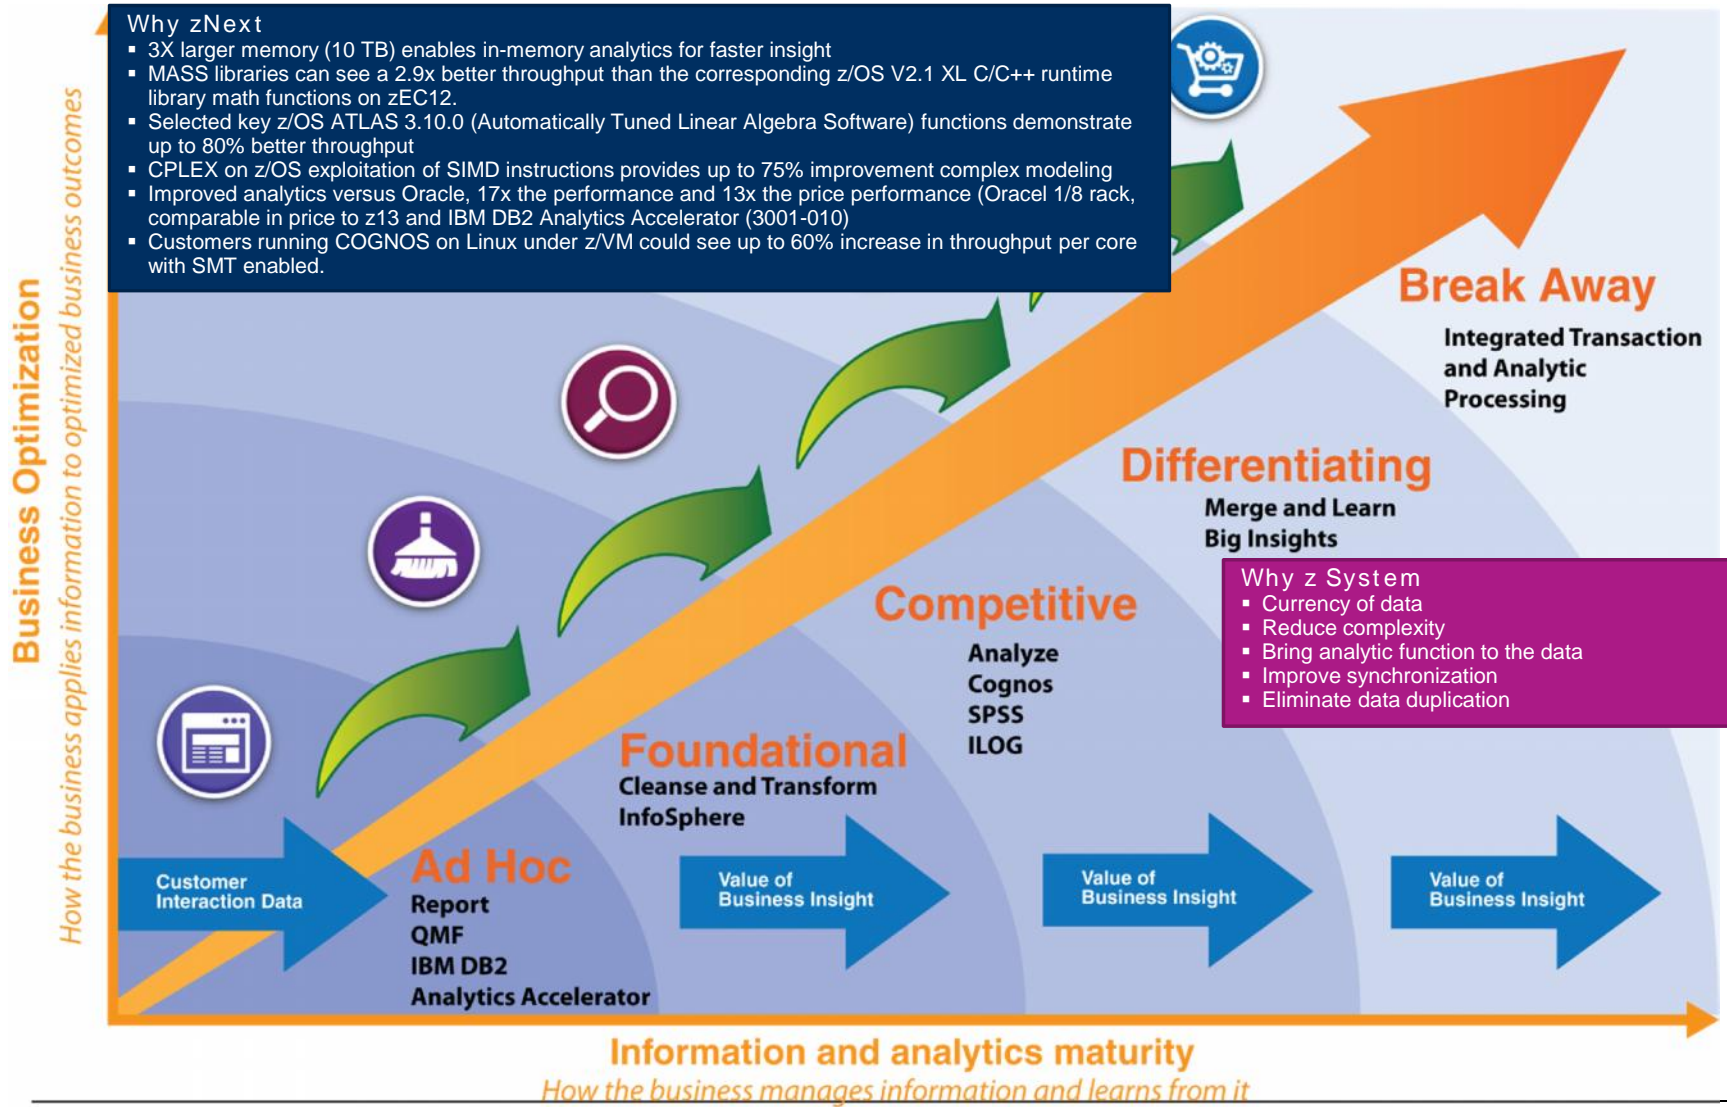
The z System strategy

Analytics as part of the flow of business - Increases richness of transactions



Best-of-class Data Lifecycle Management for:
Fighting fraud, preventing financial crimes, generating customer insights . . .

z Systems point of view



Deepen Your Knowledge of the DB2 Analytics Accelerator



IBM DB2 Analytics Accelerator

Do things you could never do before!

What is it?

The IBM DB2 Analytics Accelerator is a workload optimized, appliance add-on to DB2 for z/OS, that enables the integration of analytic insights into operational processes to drive business critical analytics and exceptional business value

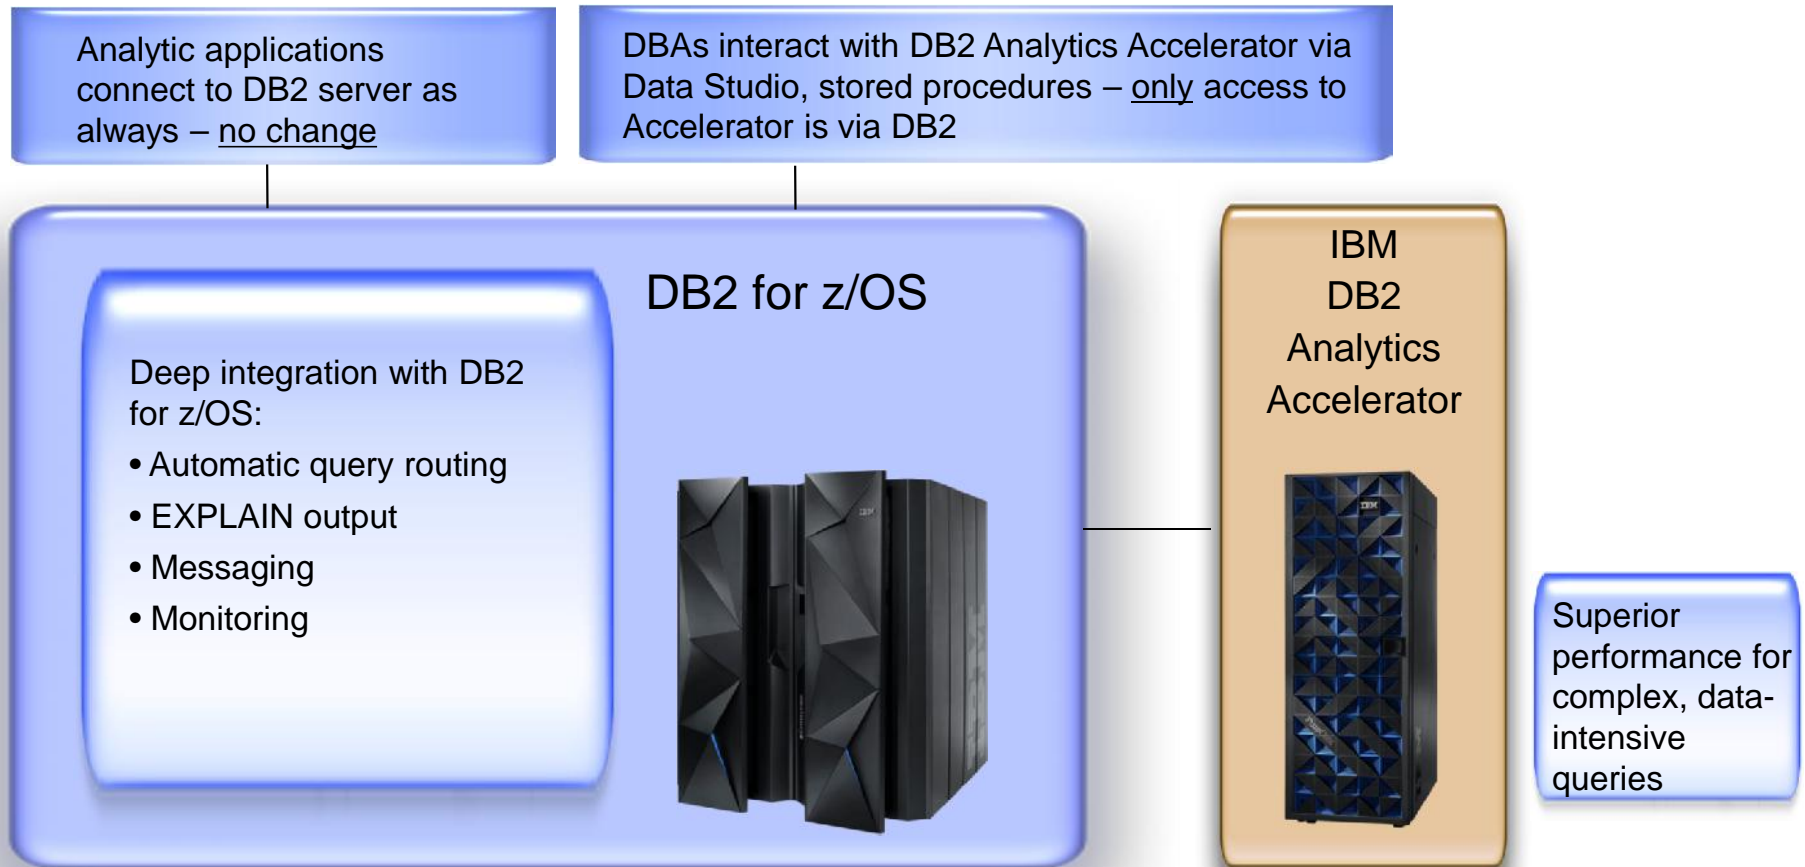
What does it do?

- Accelerates complex queries, up to 2000x faster
- Improves access to and lowers the cost of storing, managing and processing historical data
- Minimizes latency
- Reduces z Systems capacity requirements
- Improves security and governance
- Reduces operational costs and risk
- Complements existing investments

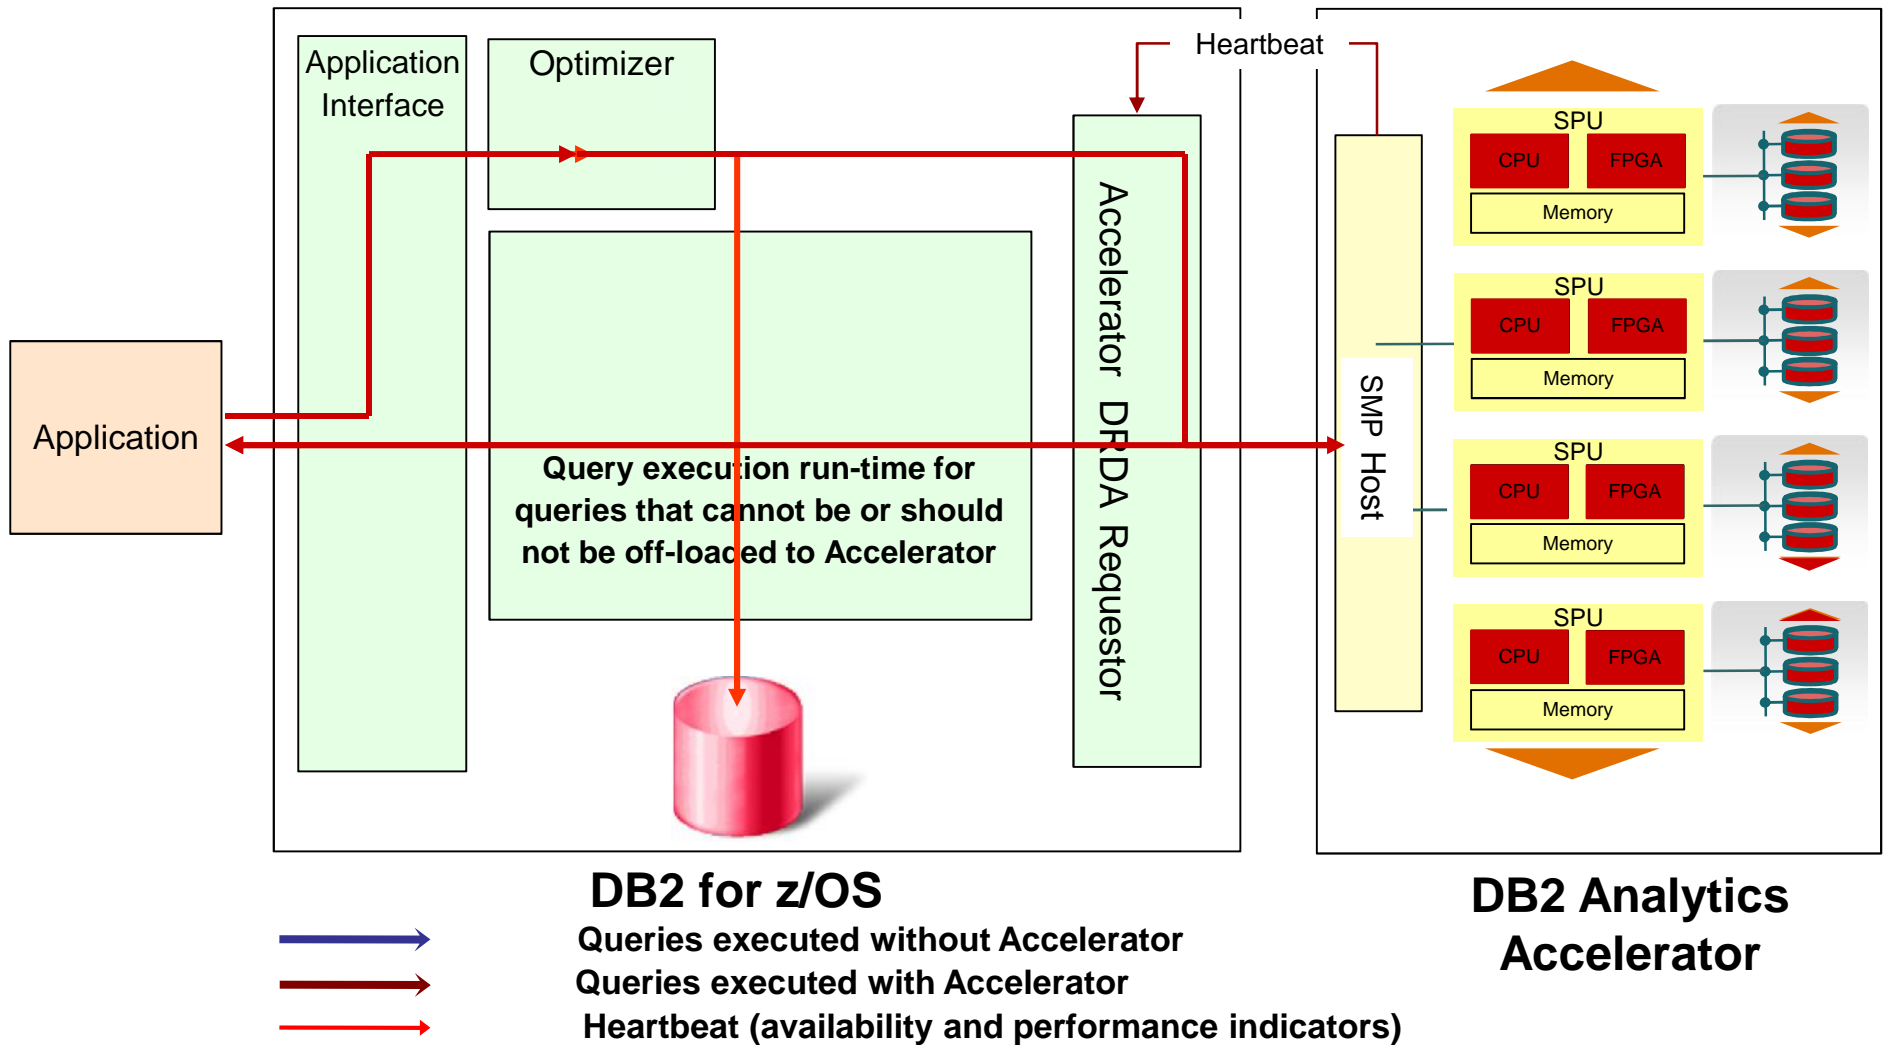


DB2 Analytics Accelerator - not “just an appliance”

An extension of the DB2 for z/OS system



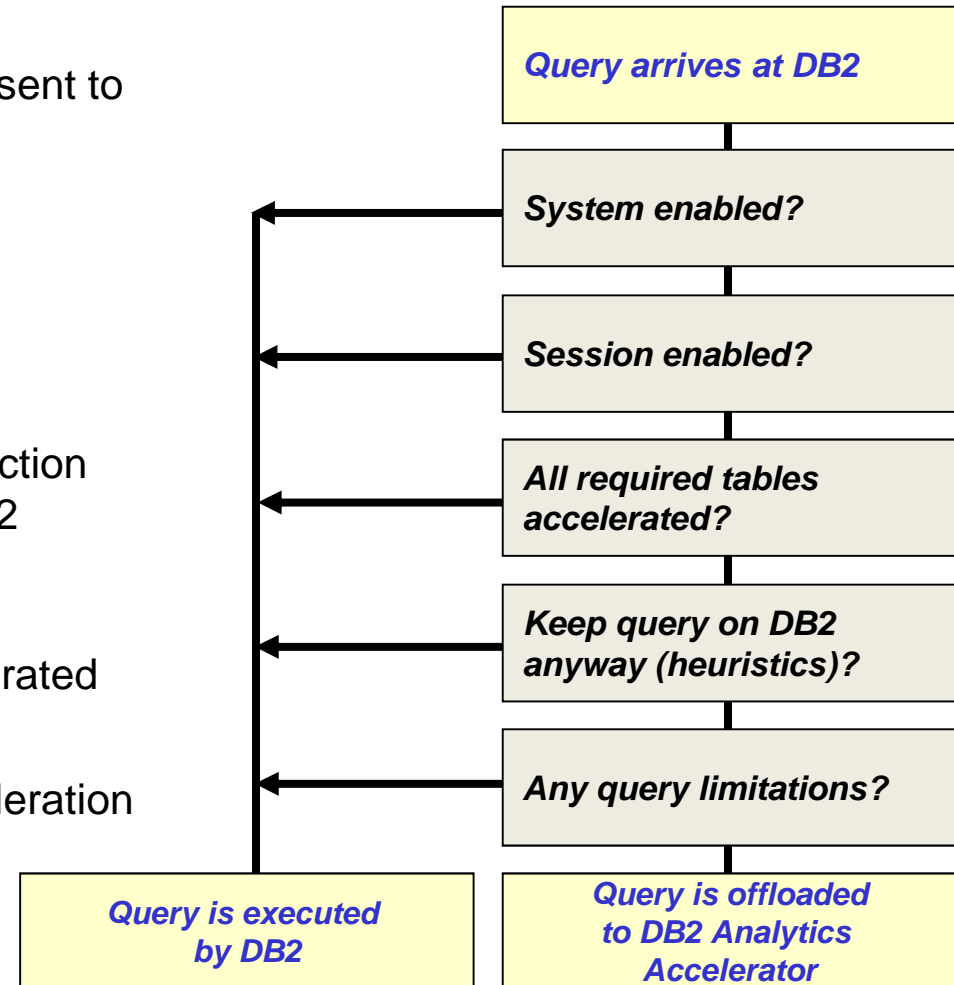
Query Execution Process Flow



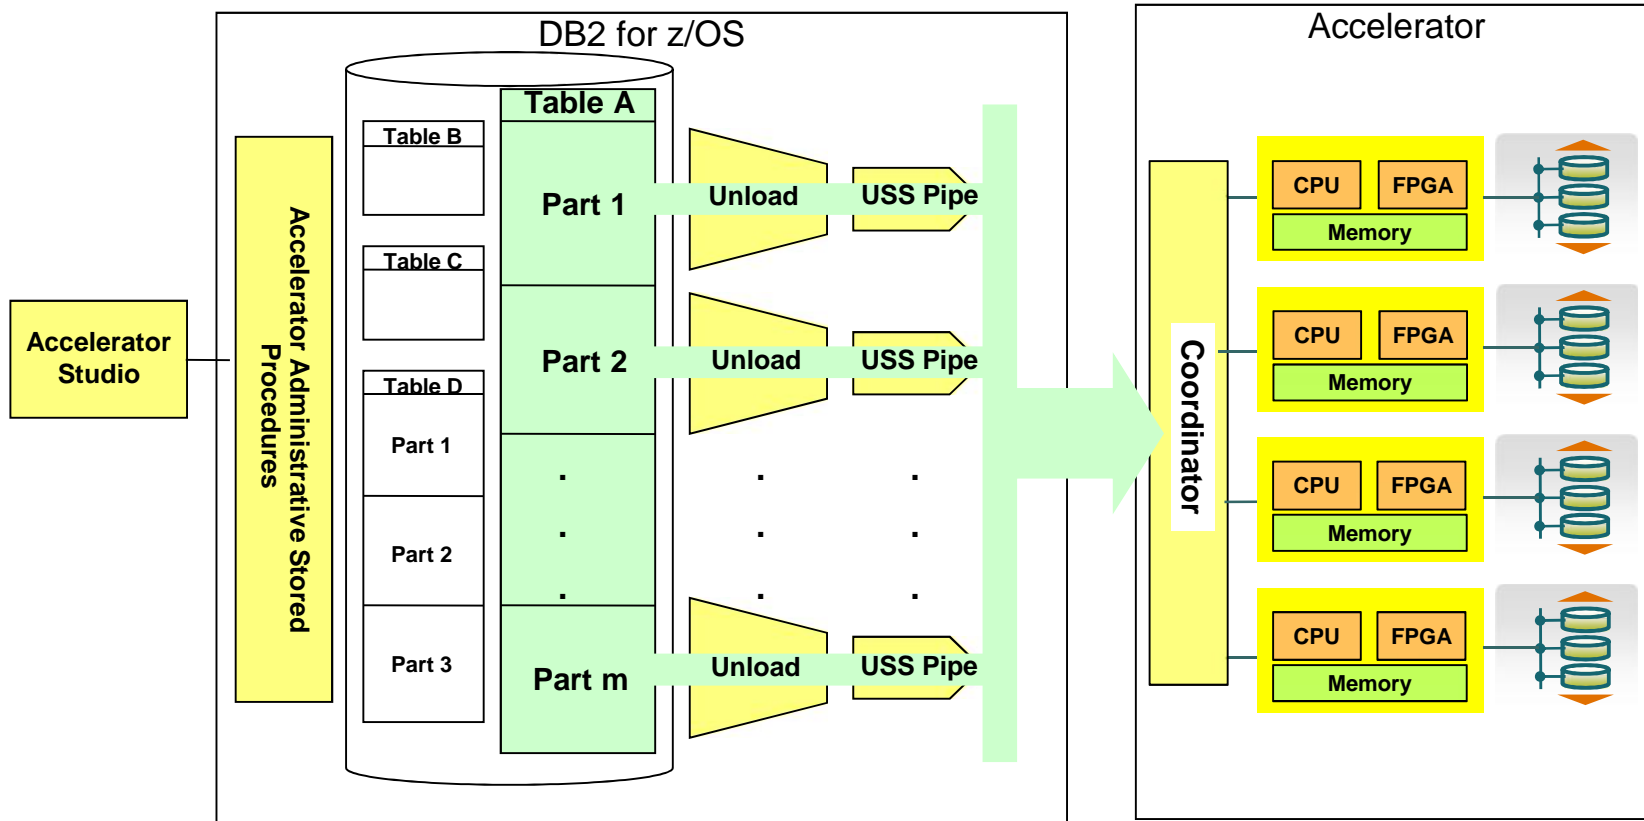
Automatic Routing Criteria

- DB2 Optimizer decides if query should be sent to accelerator
 - Special Register
 - QUERY ACCELERATION
 - NONE
 - ENABLE
 - ENABLE WITH FAILBACK
 - ELIGIBLE
 - ALL
 - Set via zPARM, Set Statement, Connection Properties / Driver, SQL Pre-Pend, DB2 PROFILES, BIND Option
- Whole query, not parts of query are accelerated
- Only read queries are considered for acceleration
- Both static and dynamic queries can be accelerated*

* Routing for static queries determined at bind time



Accelerator Data Load




- Rate can vary, depending on CPU resources, table partitioning, columns...
- Update on table partition level, concurrent queries allowed during load
- Version 2.1 & Version 3.1 unload in DB2 internal format, single translation by Accelerator

DB2 Analytics Accelerator Synchronization Options

Synchronization options	Use cases, characteristics and requirements
<p><u>Full table refresh</u></p> <p>The entire content of a database table is refreshed for Accelerator processing</p>	<ul style="list-style-type: none"> ▪ Existing ETL process replaces entire table ▪ Multiple sources or complex transformations ▪ Smaller, un-partitioned tables ▪ Reporting based on consistent snapshot
<p><u>Table partition refresh</u></p> <p>For a partitioned database table, selected partitions can be refreshed for Accelerator processing</p>	<ul style="list-style-type: none"> ▪ Optimization for partitioned warehouse tables, typically appending changes “at the end” ▪ More efficient than full table refresh for larger tables ▪ Reporting based on consistent snapshot
<p><u>Incremental Update</u></p> <p>Log-based capturing of changes and propagation to DB2 Analytics Accelerator with low latency (typically few minutes)</p>	<ul style="list-style-type: none"> ▪ Scattered updates after “bulk” load ▪ Reporting on continuously updated data (e.g., an ODS), considering most recent changes ▪ More efficient for smaller updates than full table refresh
<p><u>Accelerator Loader</u></p> <p>Separate Tool with High Value</p>	<ul style="list-style-type: none"> ▪ Image Copy Source ▪ RI without locking ▪ External Sources

Customer “A” Example



270 of the Mixed Workload Queries



Executes in DB2 returning results in seconds or sub-seconds

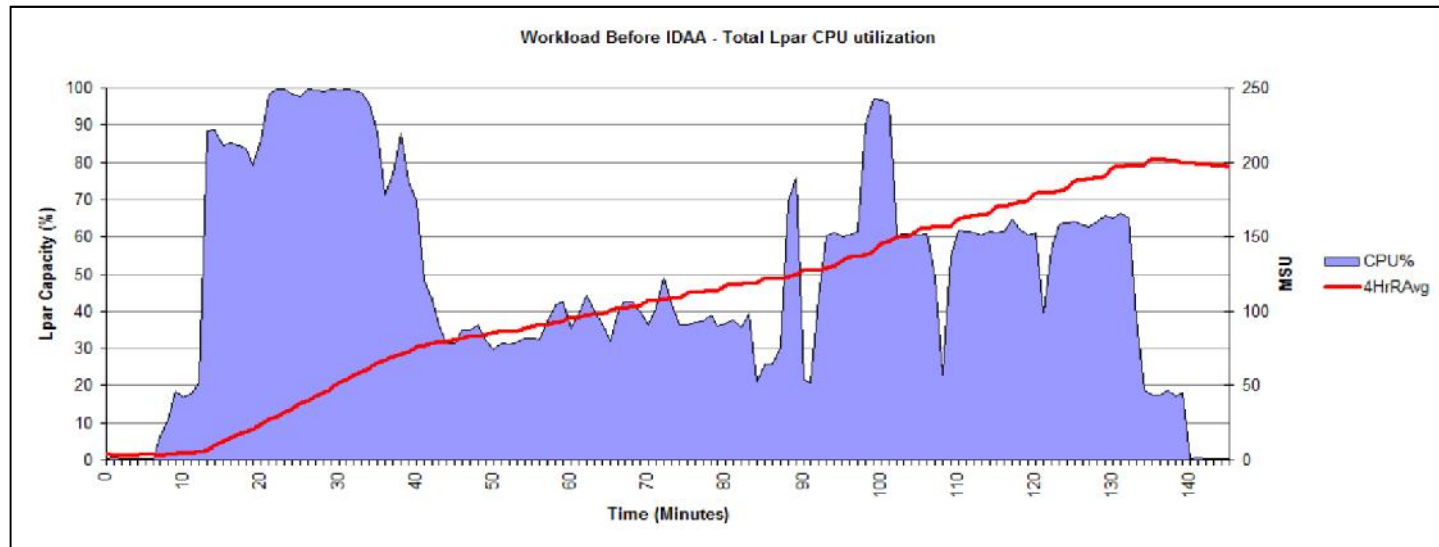
30 of the Mixed Workload Queries took minutes to hours

Query	Total Rows Reviewed	Total Rows Returned	DB2 Only		DB2 with IDAA		Times Faster
			Hours	Sec(s)	Hours	Sec(s)	
Query 1	2,813,571	853,320	2:39	9,540	0.0	5	1,908
Query 2	2,813,571	585,780	2:16	8,220	0.0	5	1,644
Query 3	8,260,214	274	1:16	4,560	0.0	6	760
Query 4	2,813,571	601,197	1:08	4,080	0.0	5	816
Query 5	3,422,765	508	0:57	4,080	0.0	70	58
Query 6	4,290,648	165	0:53	3,180	0.0	6	530
Query 7	361,521	58,236	0:51	3,120	0.0	4	780
Query 8	3,425,29	724	0:44	2,640	0.0	2	1,320
Query 9	4,130,107	137	0:42	2,520	0.1	193	13

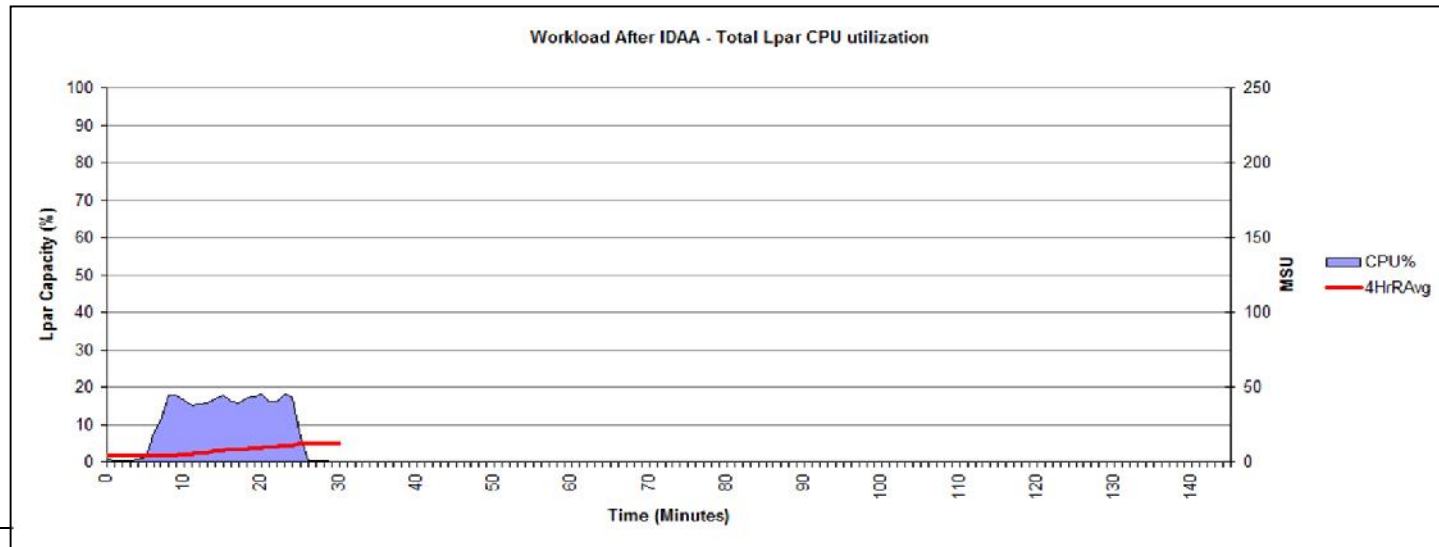
Successfully accelerated the problem queries without affecting the rest

LPAR CPU utilization comparison

*Without
accelerator*



*With
accelerator*



Fast Evolution of IBM DB2 Analytics Accelerator

- Version 1
 - IBM Smart Analytics Optimizer
 - In-memory, column-store, multi-core and SIMD algorithms
 - Discontinued and replaced by IBM DB2 Analytics Accelerator

- Version 2
 - New name: IBM DB2 Analytics Accelerator
 - Incorporates Netezza query engine
 - Preserves key V1 value propositions and adds many more

- Version 3
 - Better performance, more capacity
 - Incremental update
 - High Performance Storage Server

- Version 4
 - Much broader acceleration opportunities
 - More enterprise features



Version 4.1 at a glance – including all PTFs

More Query Acceleration	Enhanced Capabilities	Improved Transparency
<p>Static SQL</p> <p>DB2 11 ⁽²⁾</p> <p>Multi-row fetch from local applications</p> <p>EBCDIC and Unicode in the same DB2 system and Accelerator</p> <p>NOT IN and ALL predicates</p> <p>FOR BIT DATA support</p> <p>24:00:00 time and timestamp value</p> <p>MEDIAN support</p>	<p>Improved scalability and better performance of Incremental Update, more in PTF 5: Parallel Apply, Auto-Selection of keys ⁽³⁾</p> <p>Improved performance for load and large result sets</p> <p>Better access control for HPSS archived partitions</p> <p>HPSS archiving to multiple Accelerators</p> <p>Extending WLM support to local applications</p> <p>Rich system scope and statement level ⁽⁴⁾ monitoring</p> <p>Reporting prospective CPU cost and elapsed time savings</p> <p>Separation of duties for Accelerator system administration operations</p>	<p>Automatic workload balancing with multiple Accelerators</p> <p>New RTS 'last-changed-at' timestamp ⁽²⁾</p> <p>Automated NZKit installation</p> <p>Built-in Restore for HPSS</p> <p>Protection for image copies created by HPSS archiving process</p> <p>Profile controlled special registers ⁽²⁾</p> <p>Improved continuous operations for Incremental Update</p> <p>Refreshing DB2 Analytics Accelerator table without table lock even if incremental update active</p>
<p>SELECT INTO for static SQL support</p> <p>Application-period temporal table support ⁽⁴⁾</p> <p>NVL, DECODE, TRIM, full LTRIM/RTRIM support ⁽⁴⁾</p>	<p>Extended loading capabilities ⁽¹⁾</p> <p>EXPLAIN of dynamic queries against a specific Accelerator ⁽⁴⁾</p> <p>Support for N2002, N3001, N3001-001 ⁽³⁾ hardware</p> <p>Incremental Update continues replicating even for tables in AREO state¹</p> <p>In-Database transformation and multi-step processing ⁽³⁾</p> <p>Dynamic Statement Cache support for accelerated queries ⁽⁴⁾</p>	<p>Improved network bandwidth using link aggregation</p> <p>Support for Self-Encrypting disks in N3001 models</p> <p>Support for Call Home</p> <p>Archiving of DB2 transparent archive tables</p> <p>Suspending faulty tables from replication⁽³⁾</p>

(1) – delivered by a separate tool (DB2 Analytics Accelerator Loader for z/OS)
 (2) – DB2 11 only
 (3) – DB2 Analytics Accelerator Version 4.1 PTF5
 (4) – delivered by DB2 PTFs

Introducing Accelerator-only Table Type in DB2 for z/OS

Creation (DDL) and access remains through DB2 for z/OS in all cases

Non-accelerated DB2 table

- Data in DB2 only

Accelerated DB2 table

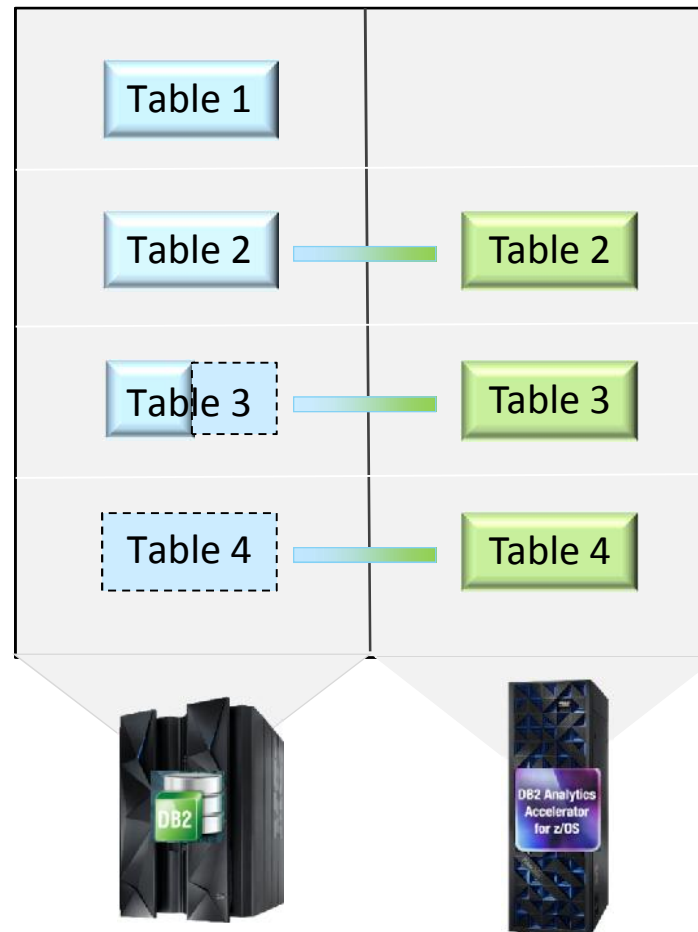
- Data in DB2 and the accelerator

Archive table / partition

- Empty read-only partition in DB2
- Partition data is in accelerator only

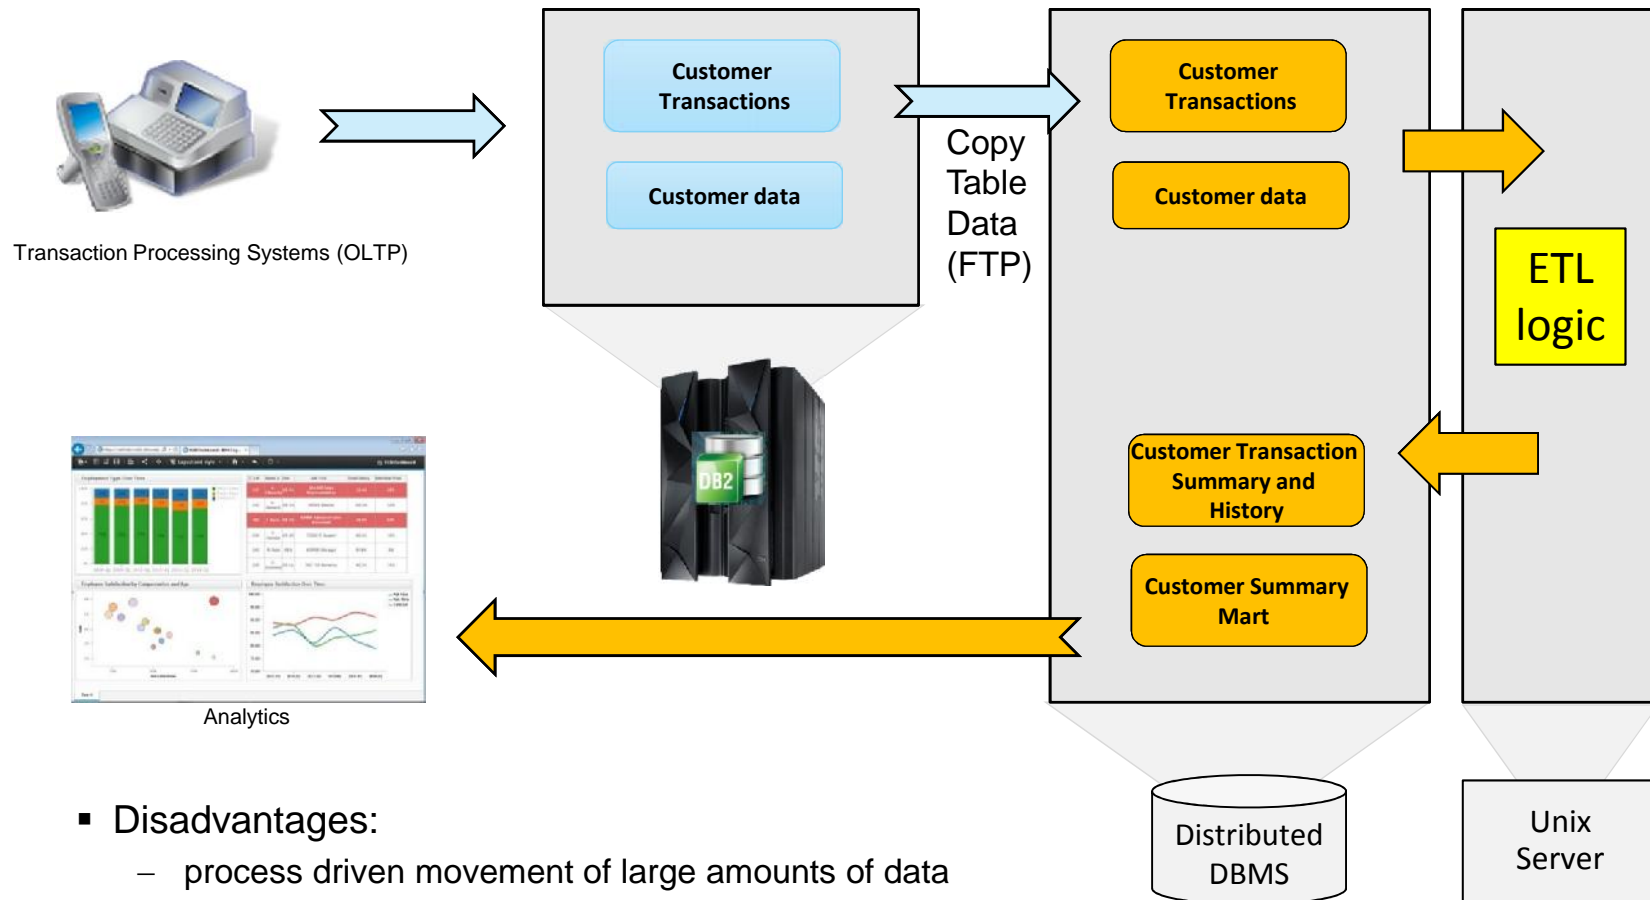
Accelerator-Only table (AOT)

- “Proxy table” in DB2
- Data is in accelerator only



In-database transformation

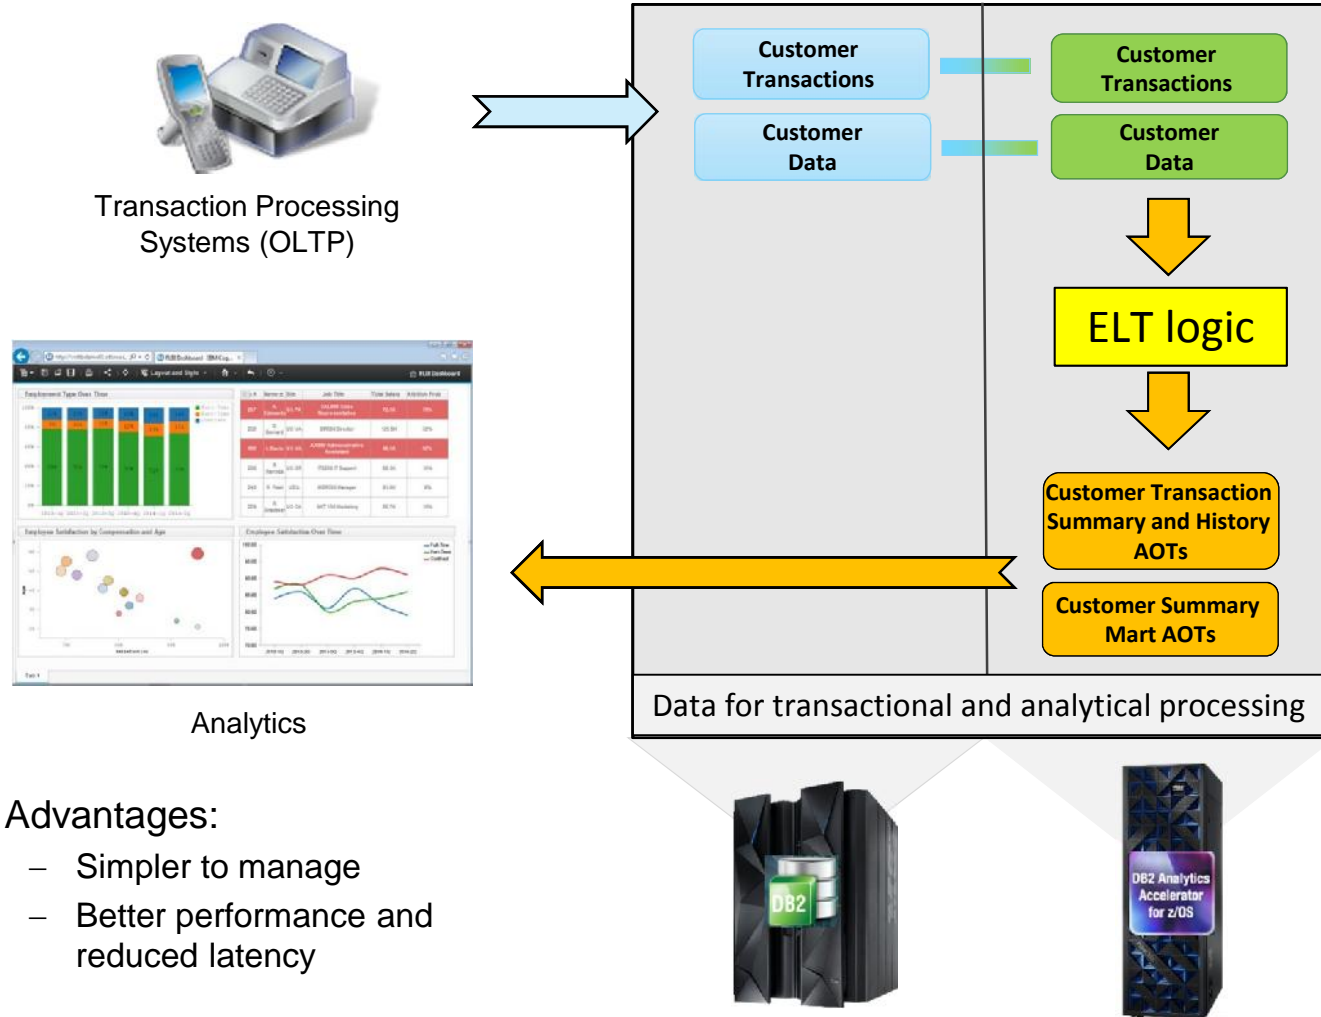
Before AOTs: ETL on a different Platform (Traditional Approach)



- **Disadvantages:**
 - process driven movement of large amounts of data
 - aged data for analytics/reporting depending on performance of data movement and transformation process

In-database transformation (cont'd.)

With AOTs: Using Accelerator-only Tables and ELT logic in the Accelerator

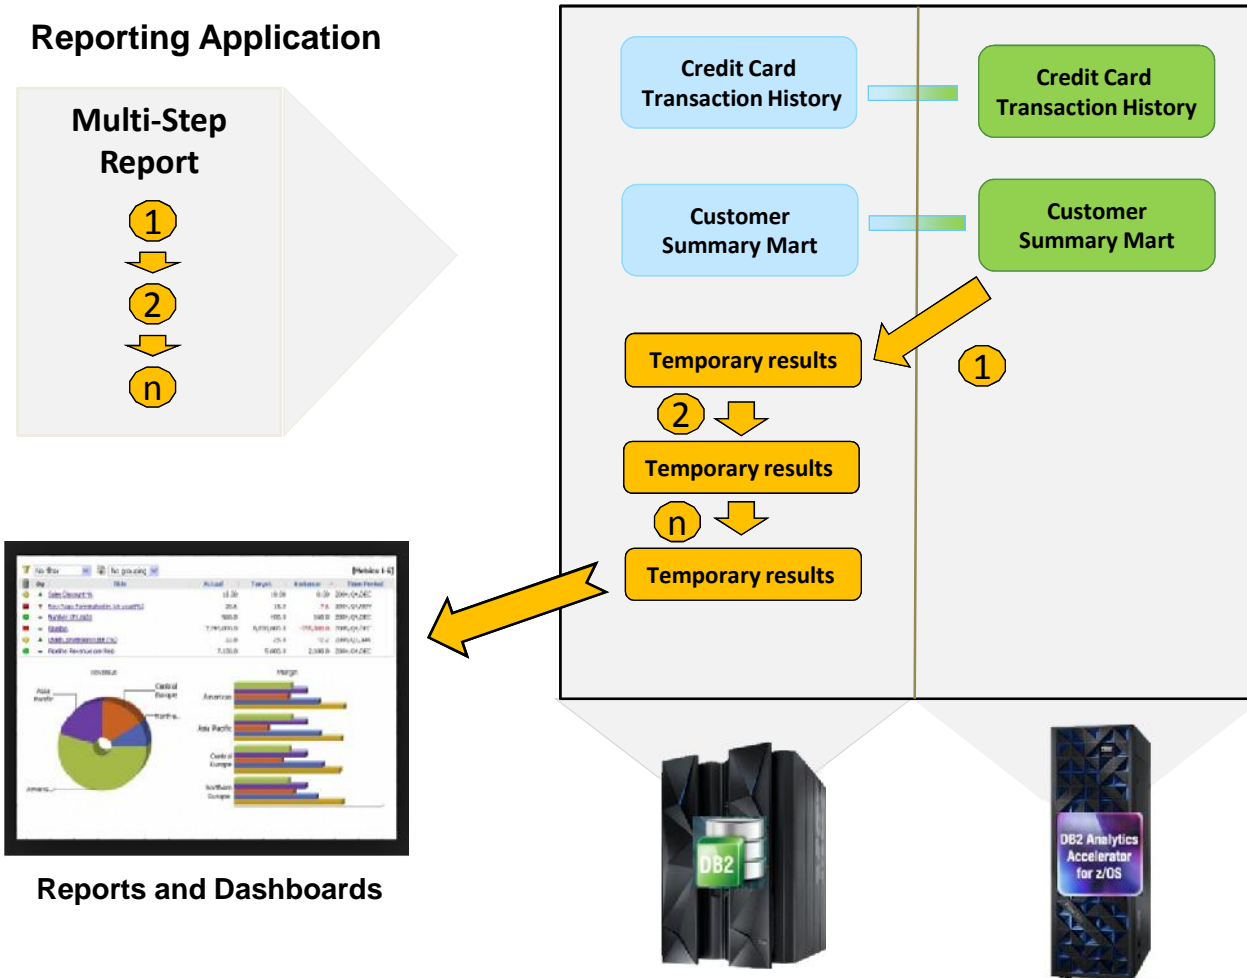


To get a backup copy of an AoT, they could be loaded into a DB2 z/OS regular or accelerated table again

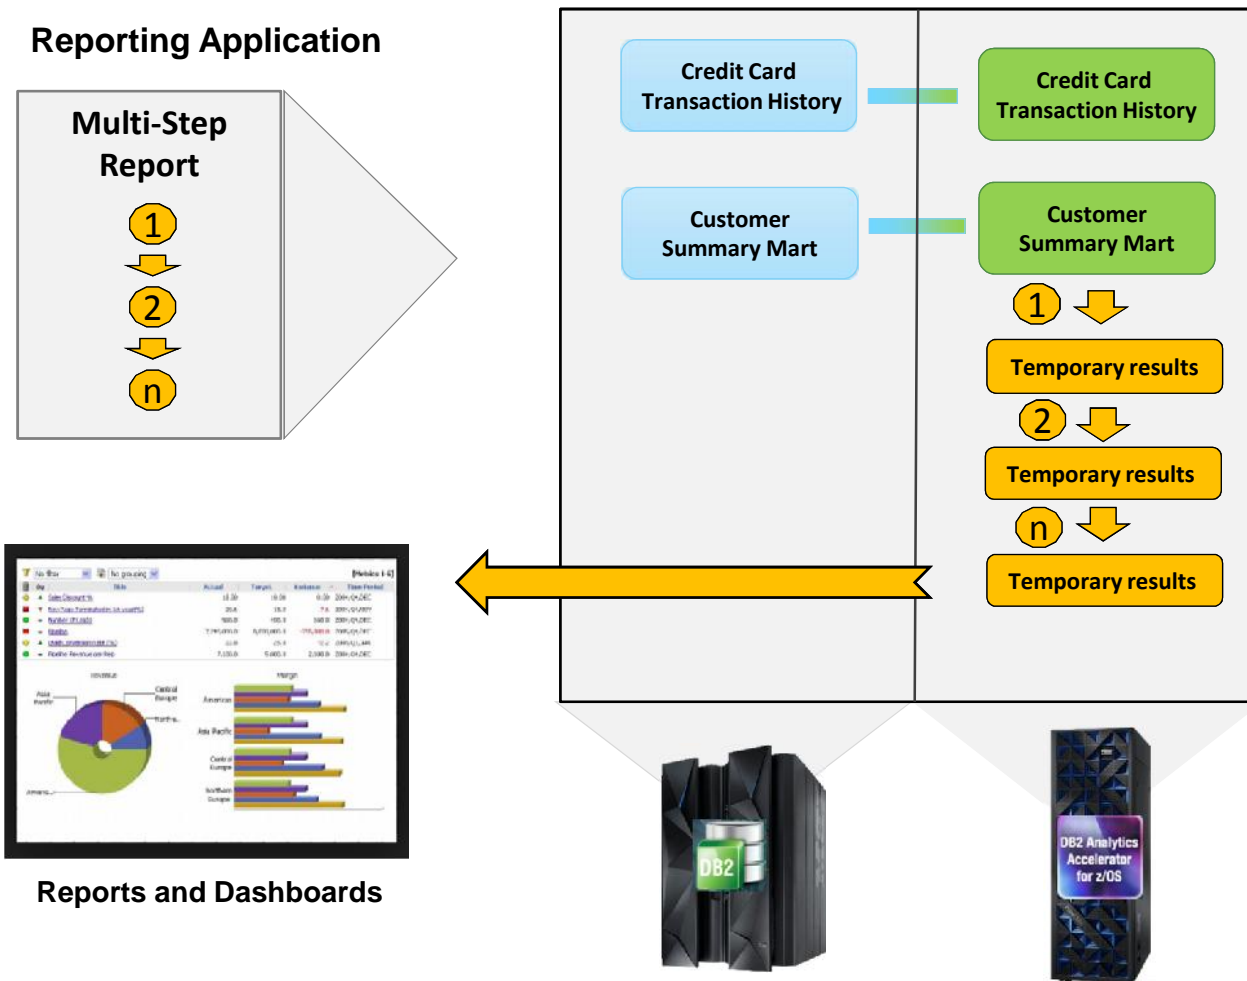
- Advantages:
 - Simpler to manage
 - Better performance and reduced latency

Multi-Step Reporting Applications with DB2 for z/OS

Before AOTs: Report processing on DB2, source data might reside on the Accelerator already



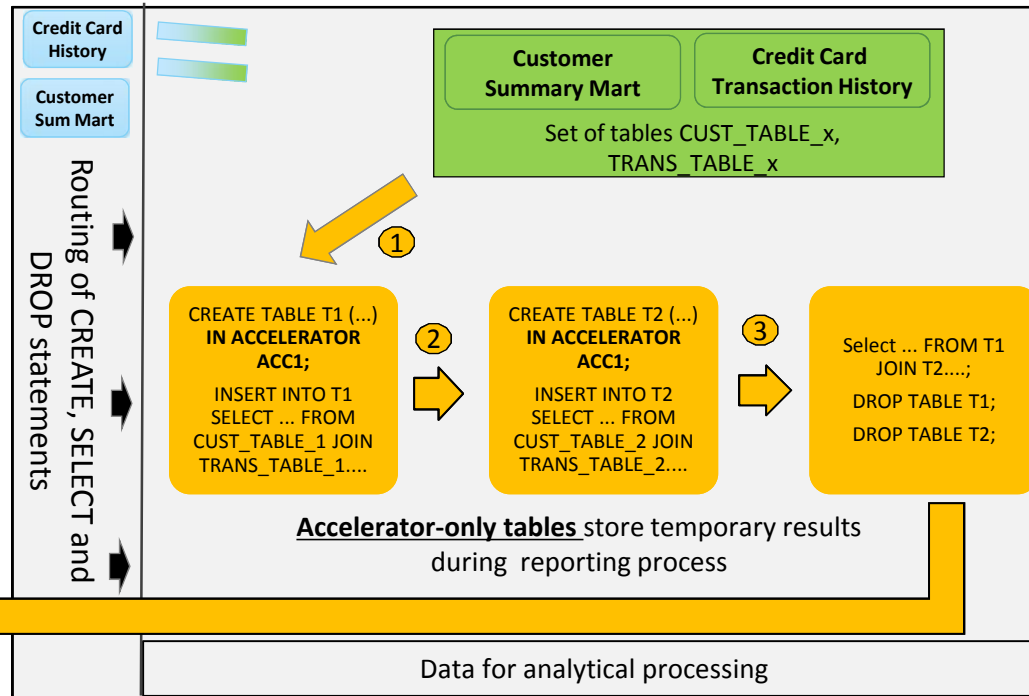
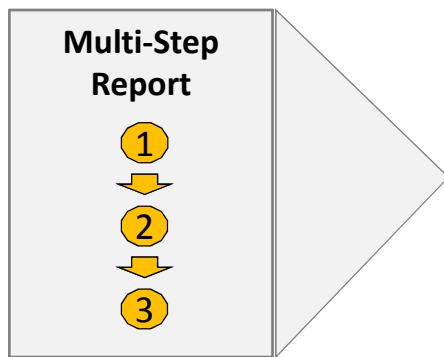
Multi-Step Reporting Applications with DB2 for z/OS (contd.)



Multi-Step Reporting Applications with DB2 for z/OS (cont'd.)

With AOTs: SQL Statement (DDL and DML) Details

Reporting Application




Reports and Dashboards



ACC1

New Use cases with Accelerator-only tables

- Adding application support for temporary objects
 - Multi-step Reporting
 - QMF
 - Campaign
- In-database transformation
 - Consolidation of ETL/ELT processing in DB2 for z/OS
 - Support for Data Stage Balanced Optimization
- Individual ad-hoc analysis
 - Data Scientist Work Area
- Integrate more data sources for analytics
 - Use DB2 Analytics Accelerator Loader to integrate with IMS data or data from other sources
 - Leverage High Performance Storage Saver data



**Determine if you can leverage the
z Systems Analytics Advantage**



Where to look for workloads

- Extending the use of operational data for business analysis, embedding operational analytics in other applications, or daily business intelligence reporting (Cognos, QMF, Business Objects)
- Long-running DB2 for z/OS queries (e.g. > 5 seconds) with at least one of the following: WHERE, GROUP BY, ORDER BY, aggregate functions
 - Queries run from a business intelligence environment that provide important business information.
 - Queries that are scheduled in batch processes overnight to not affect corporate users during the day.
- Forgotten queries that are no longer run because of performance issues.
- Analytical and ad hoc queries that may not currently run in DB2 for z/OS but could provide significant value to the end user or organization.
- Consolidating data and analytic workloads to a single secure data environment, thereby reducing organizational costs from maintenance, ETL, administration, licensing.
- Move less frequently used tables or table partitions to the Accelerator and remove the data from DB2 for z/OS.
 - Reduce the cost of storing, managing, and processing historical data while achieving improved response times for queries against the historical data.
- Analyzing non-DB2 data (e.g. IMS, VSAM, flat files) on the Accelerator.
 - Users can load non-DB2 data into the Accelerator in cases where the user would like to analyze this data.

Assess your existing workload

Analyze existing workload on your own per DB2 for z/OS member (of course, IBM offers help and can analyze SMF records for you):

To enable the “Accelerator Modeling” feature for your **DB2 10 for z/OS** (NFM or CM) subsystem, apply the following PTFs (and potential pre-req-PTFs):

- **PTF UK97693** for PM90886 for dynamic SQL assessment (included in PUT1310)
- **PTF UK98716** for PM95035 for static SQL assessment (included in PUT1311)

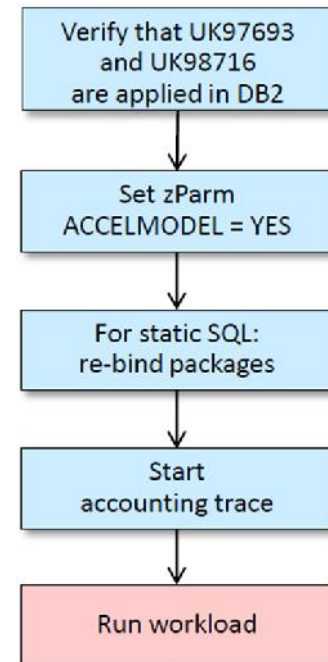
For **DB2 11 for z/OS** (NFM or CM), apply **PTF UI13488** for PM PM96478.

You may capture workload assessment data for both dynamic and static SQL at the same time (static SQL assessment requires REBIND).

You cannot run “Accelerator modeling” when query acceleration is enabled.

-START TRACE(ACCTG) CLASS(1,2) DEST(xxx) (xxx can be SMF or GTF)

It's essential to sync-up with application people to find the best timeframe for an analysis !!!



Statement Level Monitoring

- **New set of statement level monitoring counters**
 - Report SQL statement's eligibility to be routed to the Accelerator
 - Report SQL statement's execution statistics for queries that are routed to the Accelerator
 - DB2 10: Available with APAR PI23083, DB2 11: in development

Assess your existing workload

- Enabling APARs for reports to include data from accelerator modeling, including in PDB/CSV support
 - PM 99691 for OMPE 5.1.1
 - PI 08288 for OMPE 5.2.0 (DB2 10 and DB2 11)
- If you use OMPE in your environment, this new option could be another alternative for you to explore benefits with DB2 Analytics Accelerator

MEASURED/ELIG TIMES	APPL (CL1)	DB2 (CL2)
-----	-----	-----
ELAPSED TIME	4.830139	4.740227
ELIGIBLE FOR ACCEL	N/A	4.442327
CP CPU TIME	6.337894	6.336111
ELIGIBLE FOR SECP	4.990042	N/A
ELIGIBLE FOR ACCEL	N/A	6.329119
SE CPU TIME	0.000000	0.000000
ELIGIBLE FOR ACCEL	N/A	0.000000

- regular total class 1 and 2 elapsed times of the parent only
- Elapsed time spent in DB2 for SQL eligible for acceleration*. If the statements are executed in parallel, the elapsed time relates to the parent task only.
- regular total class 1 and 2 CPU times, inclusive all parallel threads
- eligible times for specialty engine processing
- The part of CPU time spent on general purpose processors for SQL eligible for acceleration *. If the statements are executed in parallel, the CPU includes the parent and all the subordinated parallel tasks.
- regular total class 1 and 2 Specialty Engine CPU times, inclusive all parallel threads
- The part of CPU time spent on specialty engine processors for SQL eligible for acceleration *. If the statements are executed in parallel, the CPU saving includes the parent and all the subordinated parallel tasks.

Disclaimer: Both the elapsed and CPU time potential saving is the projected upper limit. Namely, even if a statement is routed to the accelerator there is a DB2 processing associated with preparing and sending the statement to the accelerator and, particularly, processing associated with receiving the result set and passing it back to the application. The latter can be significant in case of very large result sets.

 * time can be significantly reduced when SQL is executed on the accelerator.

**Engage with IBM to Determine the
Best Path**



Socialization Steps

- Work with IBM to drive Business / Application Use Cases

- Companion documents available
 - One is a short summary (slides geared to application folks) that discusses the Accelerator from their point of view.
 - Second is a questionnaire regarding potential use cases that can be filled out and sent back.
 - The only thing a DBA would really have to do is add some data source examples in the questionnaire in case the recipients don't know if their application is touching DB2 z/OS data.
 - In some cases, these have been tailored by industry.

Pre-Sales based z Analytics Workshop

IBM will deliver one of several workshops regarding IBM z Analytics. These workshops are designed to help clients uncover their top business requirements related to analytics and then identify areas where the IBM z Analytics can help.

A general workshop structure would consist of:

Day 1

AM Overview of customer business, infrastructure and specific challenges

PM Overview of IBM Analytics Products

Day 2

AM IBM present to customer potential solutions based on specific business challenges

PM Open discussion on how different challenges could be met based on IBM solutions, identification of next steps and follow-ups

Education and Essential Links

- Proof of Technology Events
- Primary Product Page
 - <http://www.ibm.com/software/products/en/db2analacceforsos>
- Redbook
 - <http://www.redbooks.ibm.com/redpieces/abstracts/sg248213.html?Open&ce=ism0062&ct=swg&cmp=ibmsocial&cm=h&cr=im&ccy=us>
- Customer Testimonials
 - <https://engage.vevent.com/index.jsp?eid=556&seid=68284&code=brand>
- zBig Data & Analytics Community
 - <http://searchbusinessanalytics.techtarget.com/bigdataanalytics>



Gary Crupi, IBM Distinguished Engineer, Big Data & Analytics
gary.crupi@us.ibm.com