Session #6

The Benefits of zEnterprise Analytics

Client use cases & performance results

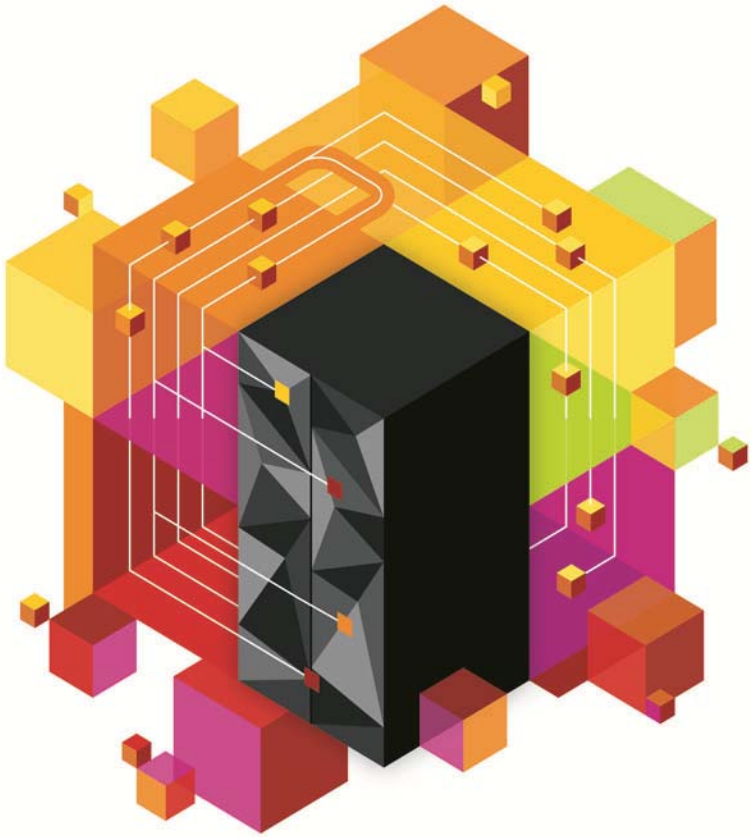

Presenter – Title

Date

Today's Discussion

- *Solutions to help maximize access to your DB2 for z/OS data for analytics*
- *Integration and optimization tips to quickly, efficiently and cost effectively get analytic solution up and running on zEnterprise*

IBM zEnterprise

an enterprise information hub providing an end-to-end, integrated foundation for modern analytics

Business Analytics
- *Business Intelligence*
- *Predictive Analytics*

Data Warehousing
- *Data Warehouses*
- *Operational Data Stores*
- *Data Marts*

Business System / OLTP

Insight Out

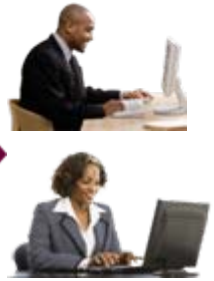

Customer Interaction

Data In

Real-Time Score/ Decision Out

Best practices for integrating & optimizing analytics into DB2 for z/OS data

Gain **30 – 45%** **Performance Improvement** with zEC12 over z196

Business Analytics
- *Business Intelligence*
- *Predictive Analytics*

Analytics Out

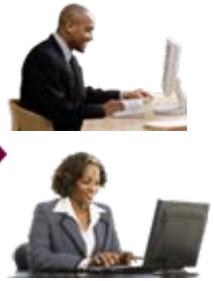

Data Warehousing
- *Data Warehouses*
- *Operational Data Stores*
- *Data Marts*

Meet SLAs & score **3000-5000+ transactions in Real-time**

Run complex queries up to **2000x Faster**

Business System / OLTP

Customer Interaction

Data In

Real-Time Score/ Decision Out

Solutions for Maximizing Access to DB2 for z/OS data

IBM DB2 Analytics Accelerator for z/OS

- **Reduces the cost of high speed analytics by blending zEnterprise and Netezza technologies to deliver, mixed workload performance for complex analytic needs**
- **Helps reduce host data warehouse storage usage by over 95%**

IBM Cognos BI for Linux on System z & z/OS

- **Reduces the cost and complexity of accessing DB2 for z/OS** data with a broad range of BI capabilities: reporting, analysis, dashboards, real-time monitoring, all on a single infrastructure. Author, share and analyze reports that draw on DB2 for z/OS data for better business decisions.

IBM SPSS Modeler for Linux on System z

- **Increases the speed and accuracy of scoring in real time by imbedding the scoring algorithm in DB2 for z/OS and running it directly within the transactional application**

IBM DB2 Analytics Accelerator

Blends zEnterprise and Netezza technologies to deliver, mixed workload performance for complex analytic needs

Fast

Complex queries run up to 2000x faster while retaining single record lookup speed

Cost Saving

Eliminate costly query tuning while offloading complex query processing

Appliance

No applications to change, just plug it in, load the data, and gain the value

Reduces the Cost of High Speed Analytics

- Choice of historical data location – High Performance Storage Saver
- Real time analytics – Incremental Update
- Faster data refresh – Unload Lite
- New capacity – Full range of Netezza models supported
- New queries

Typical Customer Example:

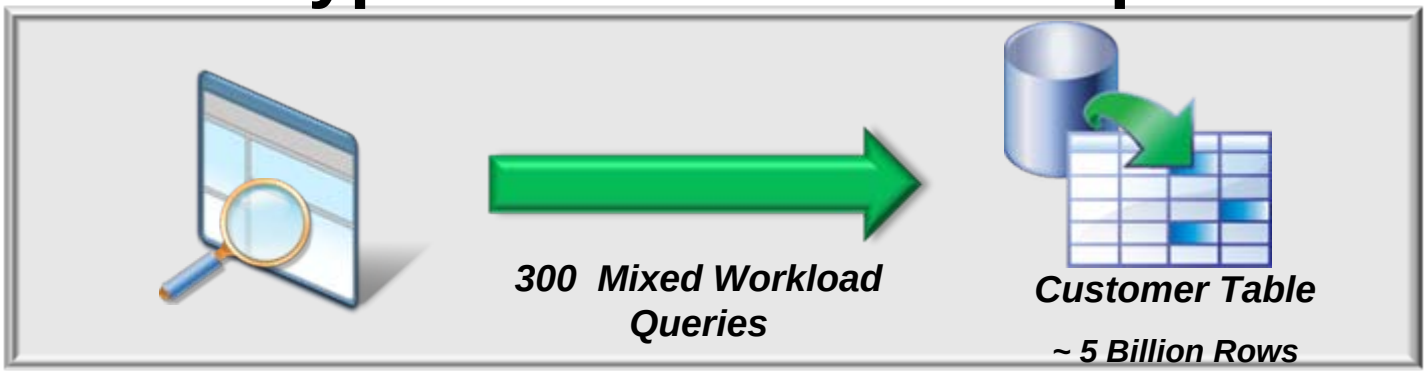

270 of the Mixed Workload Queries

Executes in DB2 returning results in seconds or sub-seconds

30 of the Mixed Workload Queries took minutes to hours

			DB2 Only		DB2 with IDAA		Times Faster
Query	Total Rows Reviewed	Total Rows Returned	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

Typical Experience - Fast Time to Value

- **IBM DB2 Analytics Accelerator**
  - ➔ **Production ready - 1 person, 2 days**
- **Table Acceleration Setup … 2 Hours**
  - DB2 “Add Accelerator”
  - Choose a Table for “Acceleration”
  - Load the Table (DB2 copy to Netezza)
  - Knowledge Transfer
  - Query Comparisons
- **Initial Load Performance …**
  - ➔**400 GB “Loaded” in 29 Min**
    570 million rows (Loads of 800GB to 1.3TB/Hr)
- **Actual Query Acceleration … 1908x faster**
  - ➔**2 Hours 39 Minutes to 5 Seconds**
- **CPU Utilization Reduction**
  - ➔**35% to ~0%**

Actual customer results, October 2011

Moving the Data Warehouse back to zEnterprise
Performance Based Quotable Quotes…

- **During first query submissions… "The accelerated version of that query just finished in 8 seconds, the DB2 version is still running"… and it ran for 27 more minutes!**

- **"That one used to run in 50 minutes, now it runs in 47 seconds…. wait a minute, last time it ran in 34 seconds… oh, who cares when you are used to 50 minutes!"**
  - The difference ended up being a heavy LOAD in progress.
  - This "tens of minutes down to seconds" theme carried forward.

- **Asked the user to remove the filters in Business Objects, the response was "Are you sure you want me to do that?" as normally they would wait forever. When the screen came back at the "speed of light", the user was "Ecstatic".**

‘Wicked Fast’ Performance for Complex DB2 Queries

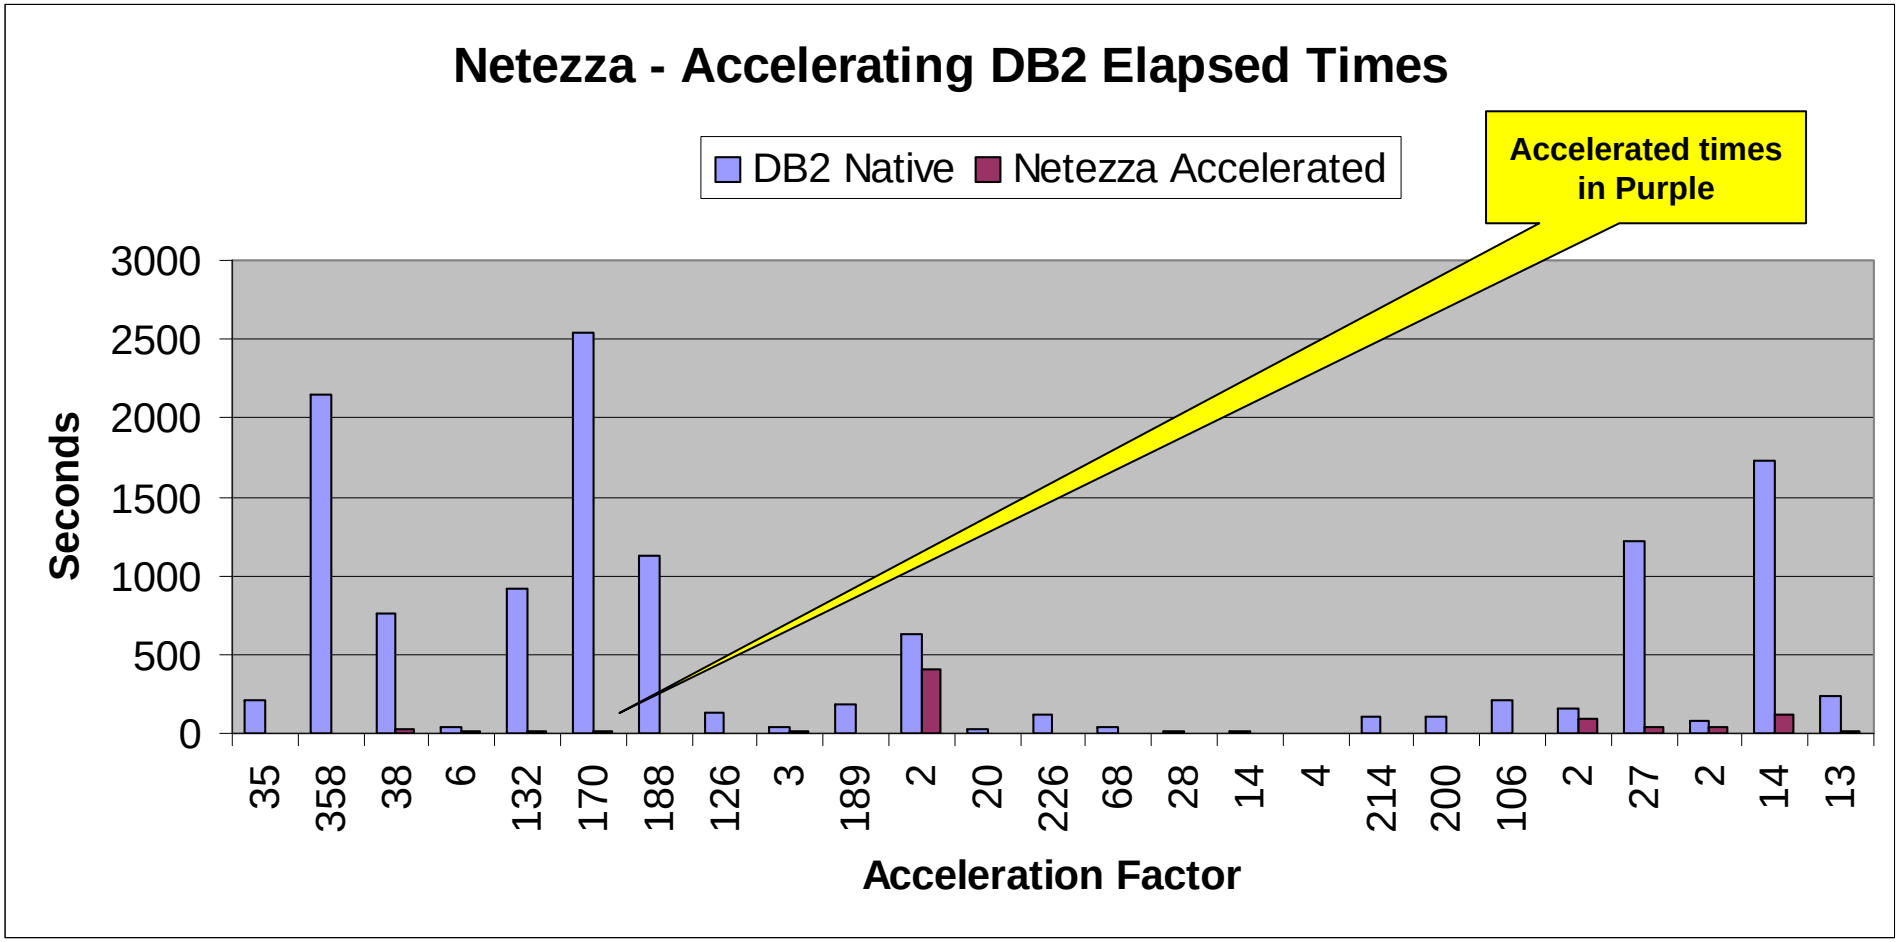

Query Details

Several Ran Sub-Second

Up to 358x Faster!

Accelerated EXEC_DATE	Without Acceleration (seconds rounded)	Without Acceleration CPU (seconds rounded)	Accelerated (seconds rounded)	X Factor	ROWS_RETURNED	RESULT SIZE	STMT_TEXT
5/9/12 12:27 PM	210	7	6	35	26787	628K	T3.CNTR_NMBR, T3.MBR_ID,
5/9/12 12:43 PM	2146	261	6	358	3868	60.4K	T1.PRFX_CD,
5/9/12 12:45 PM	755	33	20	38	624	30.9K	T1.PRFX_CD,
5/9/12 12:46 PM	40	14	7	6	72627	2.98M	T1.CNTR_NMBR, T1.MBR_ID,
5/9/12 12:47 PM	921	186	7	132	4799	136K	PRFX_CD, CNTR_NMBR,
5/9/12 12:48 PM	2543	1695	15	170	3826	299K	T1.MBR_TRNS_EFCTV_DT,
5/9/12 1:02 PM	1128	536	6	188	122948	2.81MB	CNTR_NMBR,
5/9/12 12:56 PM	126	43	1	126	0	0 B	S.CNTR_NMBR, S.MBR_ID,
5/16/12 12:00 AM	2	0	not accel (not expected to)				
5/10/12 8:15 AM	34	13	11	3	528701	9.38 MB	TBL1.PRFX_CD,
5/10/12 8:37 AM	0	0	not accel (not expected to)				A.FR_MKT_SBSEG_CD AS
5/10/12 8:48 AM	189	24	1	189	131508	1.63 MB	AS CL_MKT_SEG_CD,
5/10/12 9:31 AM	623	418	409	2	18904956	631 MB	K_AFLT_ID, AFLT_ID AS
5/10/12 9:47 AM	20	8	1	20	131508	1.38 MB	K_CNTR_ID, CNTR_NMBR
5/10/12 9:51 AM	0	0	not accel (not expected to)		299		A.K_BS_RPT_UNT_ID AS
5/10/12 12:54 PM	113		1	226	1	33B	'FS AAE NONCNTR',

Avoided CPU Consumption

Avoided Redirecting

Existing Data Warehouse
Performance Based Quotable Quotes…

- **Sample query run... native DB2 800 seconds elapsed time, DB2 Analytics Accelerator ran it in 6s 610ms.**
- **You are like a dealer… Now I expect everything to be this fast**
- **It is really impressive… queries that did not finish in DB2 can now be run**
- **178x faster is the current leader (among those that finished in DB2)**

Service Provider Query Results

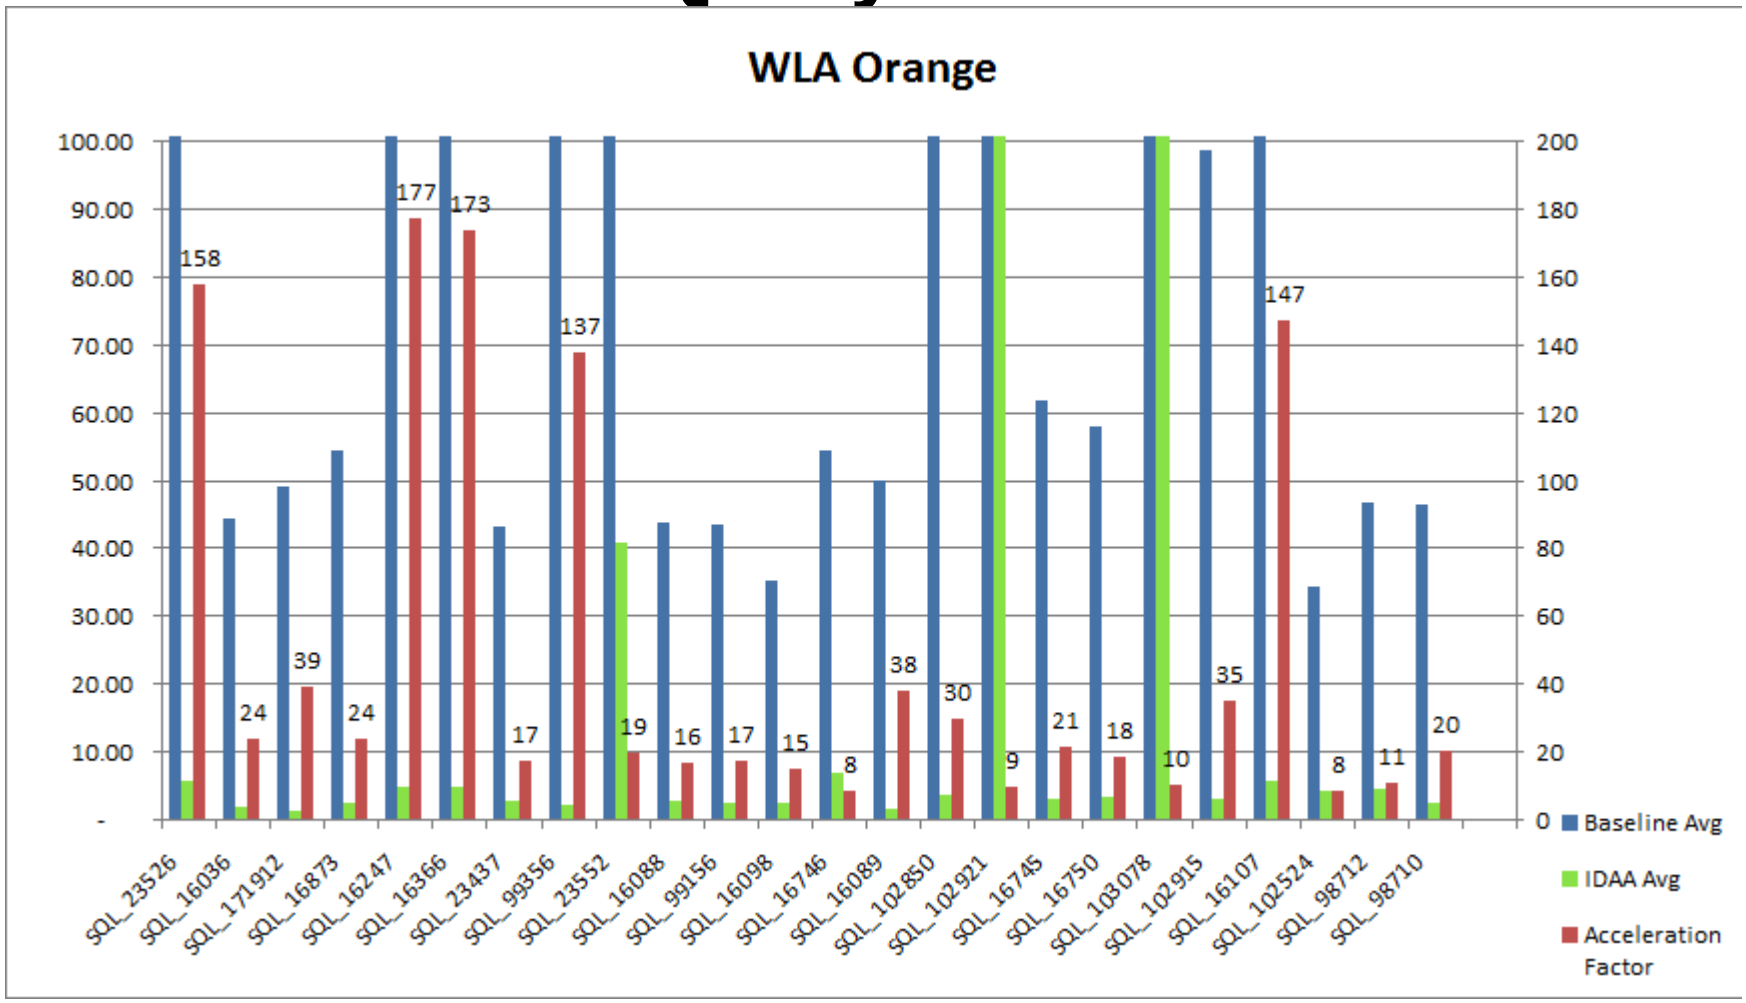

Extending Netezza Technology to zEnterprise Data and Processes

- **Bring Massively Parallel Processing (MPP) technology to System z processes to accelerate queries and processes**
  - If it already runs on z, or all the data is sourced by z
  - Leverage core infrastructure, processes and people
- **Reduce cost of running existing processes**
- **Run queries that hit the Resource Limit Facity (RLF) limit before**

Operational Workload Assessment

Query Summary

	Total	DB2 natively	With potential	Uncertain	W/o potential
Queries	11804	10896 (92%)	908 (8%)	0 (0%)	0 (0%)
Elapsed Time (s) [1*]	35813.03	11127.10 (31%)	24685.92 (69%)	0.00 (0%)	0.00 (0%)
Elapsed Time (s)	42416.47	15134.51 (36%)	27281.96 (64%)	0.00 (0%)	0.00 (0%)
CPU Time (s) [1*]	4494.65	955.70 (21%)	3538.95 (79%)	0.00 (0%)	0.00 (0%)
CPU Time (s)	5631.97	1260.64 (22%)	4371.32 (78%)	0.00 (0%)	0.00 (0%)

[1*] - Considers each query only once

Ignore non-select statements

Query classification	Unique Queries	Executed Queries	CPU %	CPU eligible %
Total	11804	941528	100%	78%
> 60 min elapsed time (info)	0	0	0%	0%
10-60 min elapsed time (info)	2	2	2%	0%
1-10 min elapsed time (info)	183	196	66%	51%
< 1 min elapsed time (info)	11619	941330	32%	27%

Reason breakdown for 0% queries with no potential and 0% of elapsed time with no potential

Reason	# Queries	% Queries	% Elapsed Time

Start trace time	End explain time	Min time stmt cached	Max time stmt cached
Apr 22, 2012 5:27 AM	May 8, 2012 12:36 PM	Apr 22, 2012 5:48 AM	May 8, 2012 12:36 PM

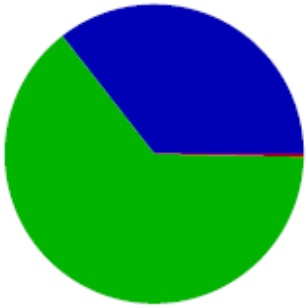

Elapsed time for best DB2 native processing
Elapsed time with acceleration potential
Elapsed time with uncertain potential
Elapsed time without acceleration potential

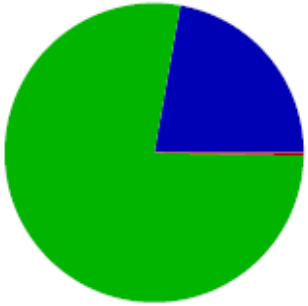

Query CPU time for best DB2 native processing
Query CPU time with acceleration potential
Query CPU time with uncertain potential
Query CPU time without acceleration potential

The experience so far…

- **Up and running in 1 day, quick start training completed**

- **1300+ tables loaded**

- **First query… in DB2, it ran for 11 minutes, 31 seconds 41 milliseconds. In the DB2 Analytics Accelerator, it ran for 1 second and 576 milliseconds.**
  - Customer really wants to shave off that 576 milliseconds.... we are working on it.

- **Typical query…**
  - DB2 Native Run - Elapsed Time 17 minutes, Normalized CPU Time 1 hour 33 minutes, **Cost $166.00**
  - DB2 Analytics Accelerator Run - Elapsed Time 59 seconds, Normalized CPU Time 10 minutes, **Cost $12.50**
  - Quote of the day… **"Just think, these run nightly and there are hundreds of them."**

Cognos BI v10.2 for Linux and z/OS on zEnterprise

Improved

User Experience

Enable users to focus on what's important

Improved

Performance & Scale

Meeting the demands of growing data volumes

* Not supported z/OS

Workloads tested

Scalability testing with Operational Analytics applications

- OpBI
- Saved Report

Only classic Cognos BI functions utilized

Functional certification testing

- 13 workloads

Workload - Operational BI (OpBI)

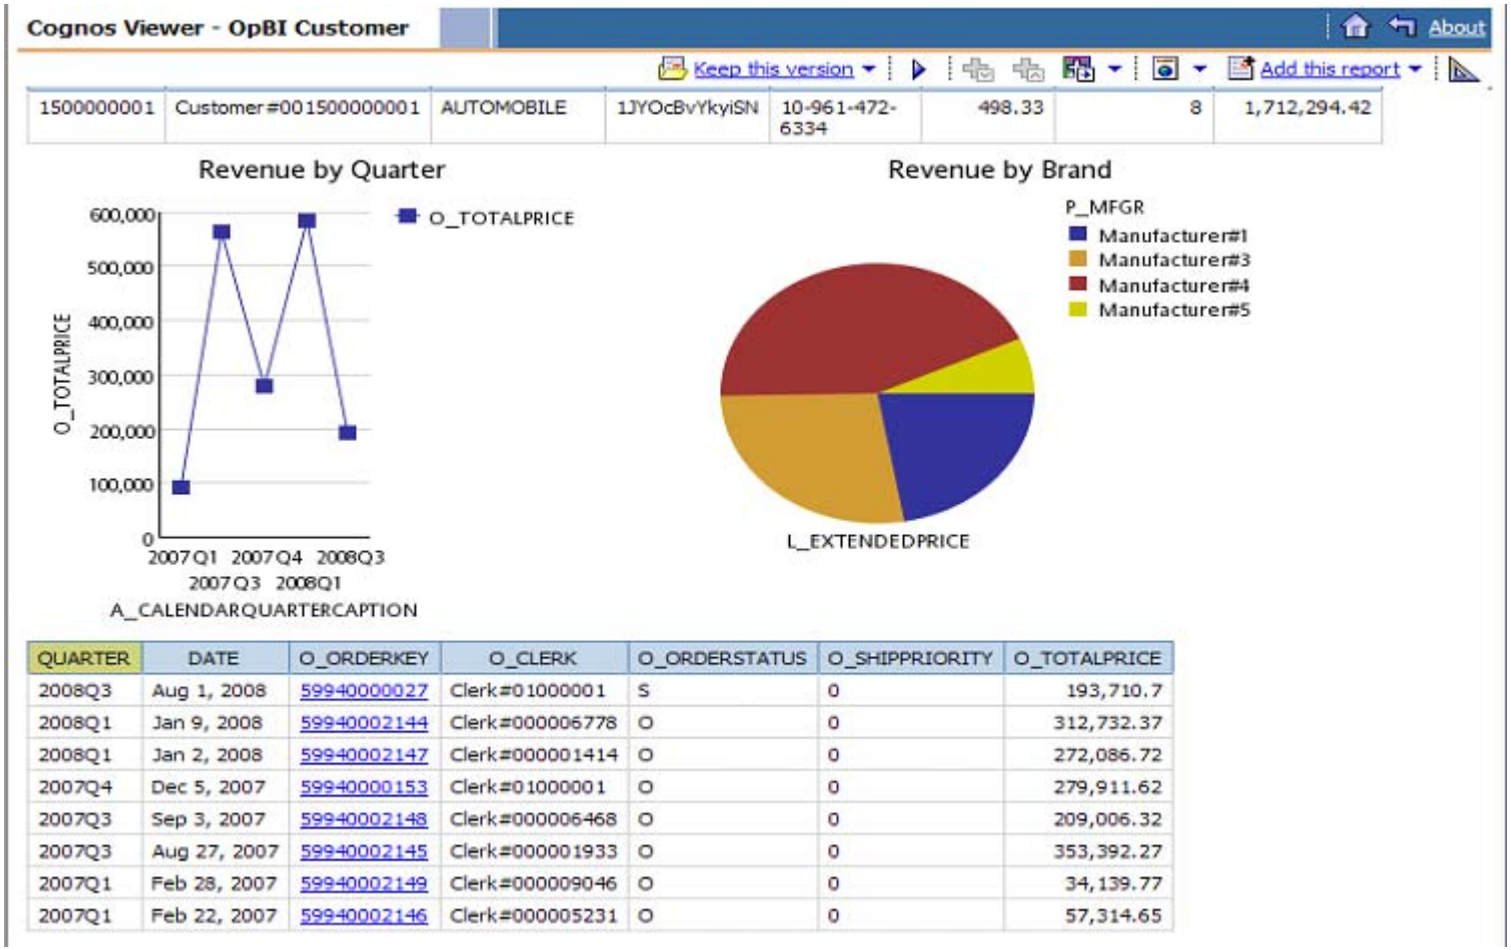

QUARTER	DATE	O_ORDERKEY	O_CLERK	O_ORDERSTATUS	O_SHIPPRIORITY	O_TOTALPRICE
2008Q3	Aug 1, 2008	59940000027	Clerk#01000001	S	0	193,710.7
2008Q1	Jan 9, 2008	59940002144	Clerk#000006778	O	0	312,732.37
2008Q1	Jan 2, 2008	59940002147	Clerk#000001414	O	0	272,086.72
2007Q4	Dec 5, 2007	59940000153	Clerk#01000001	O	0	279,911.62
2007Q3	Sep 3, 2007	59940002148	Clerk#000006468	O	0	209,006.32
2007Q3	Aug 27, 2007	59940002145	Clerk#000001933	O	0	353,392.27
2007Q1	Feb 28, 2007	59940002149	Clerk#000009046	O	0	34,139.77
2007Q1	Feb 22, 2007	59940002146	Clerk#000005231	O	0	57,314.65

Workload -- Saved Report View

Cognos Viewer - Part vs Supplier Info

Keep this version | Add this report | Watch new versions | Add comments

Not a Real Company

Part & Supplier Info

Order ID	Line #	Customer ID	Part ID	Supplier ID	Quantity	Extended Price	Discount %	Tax %	Return Flag	Line Status	Ship Date	Date	Receipt Date	Ship Instructions	Ship Mode	Discounted Price	Total Tax	Extended Revenue	Extended Supply Cost	Gross
599400000027	1	1500000001	1480645	100000001	33	$53,643.81	0.10	0.03	N	S	Oct 23, 2008	Oct 1, 2008	Oct 25, 2008	COLLECT COD	REG AIR	$48,279.43	$1,448.38	$49,727.81	$10,502.58	$37,776.85

Part ID	Part Name	Manufacture	Brand	Type	Size	Container	Retail Price
1480645	cream sandy burnished powder plum	Manufacturer #3	Brand#34	ECONOMY PLATED TIN	27	SM JAR	$1,625.57

Action	Part ID	Available Qty	Supply Cost	Supplier ID	Supplier Name	City	Nation	Phone #	Account Balance
Re-Route	1480645	0	$318.26	100000001	Supplier#000100000001	ANNABA	ALGERIA	10-309-927-9636	$5,576.87
Re-Route	1480645	3,587	$318.26	76480646	Supplier#076480646	DESE	ETHIOPIA	15-620-985-6423	$7,283.30
Re-Route	1480645	396	$800.46	1480646	Supplier#001480646	REIMS	FRANCE	16-554-633-3447	$4,669.43
Re-Route	1480645	798	$932.42	26480646	Supplier#026480646	BOMBAY	INDIA	18-292-578-9659	$6,616.01

Available Qty
3,587
500
250
0

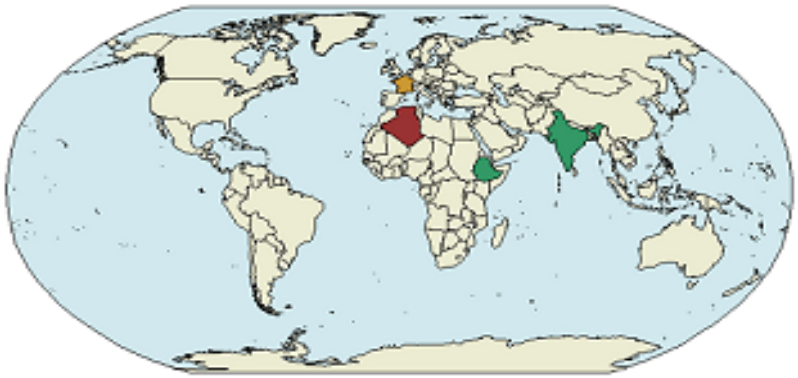

Proof Positive: Superior Scalability & Performance
Linear Scalability with Very High Demand

Pre-processed Reports

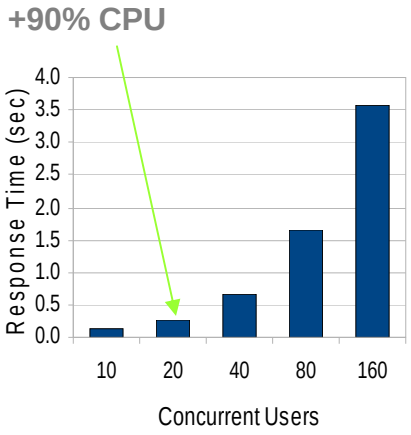

Dynamic Reports

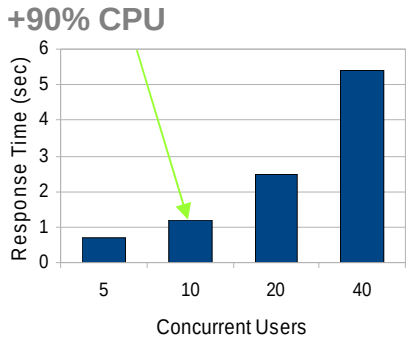

Cognos BI for Linux on System z

- Linear scalability of response time as user community grows
- Confidently accommodate the peaks with predictable service

Service the business at peak volumes without additional investment

Proof Positive: Rapid Deployment & Expansion
Up and running in minutes to meet SLAs

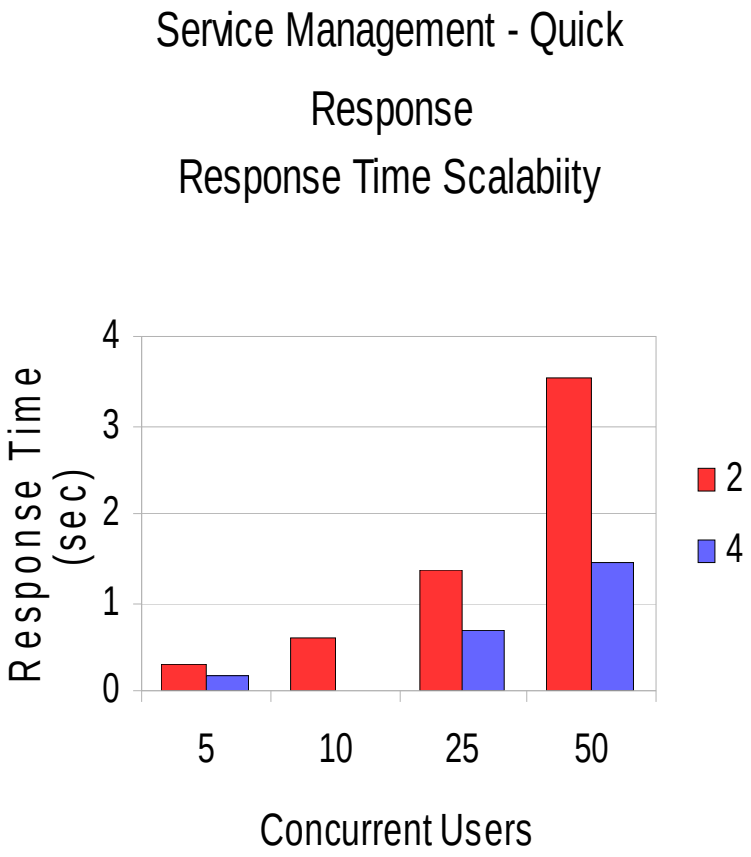

Cognos BI for Linux on System z

- Proven dynamic CPU LPAR and OS resource sharing (over 40 years of automation technology)
- Proven dynamic on demand virtualization and manual or new, automated addition of hardware resources to virtual servers
- Proven monitoring and accounting tools

Allocate hardware resources in seconds or minutes to meet your SLAs

New 5.5 GHz 6-Core Processor Chip

Large caches to optimize data serving

Second generation OOO design

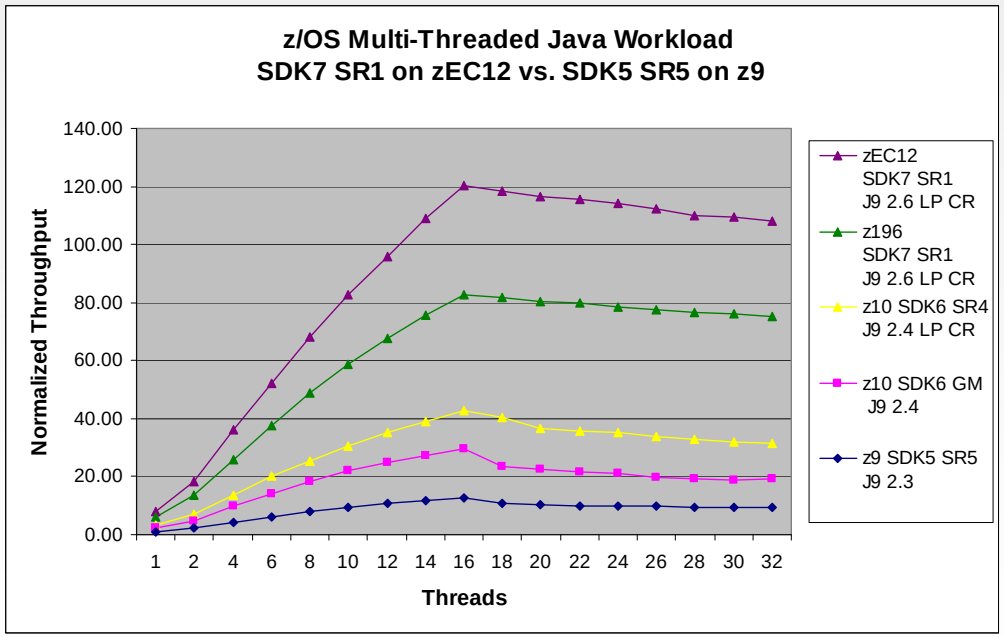

Up-to 45% improvement in throughput amongst Java workloads measured with zEC12

Multi-threaded workload shows ~11x aggregate hardware and software improvement comparing Java5SR5 on z9 to Java7SR1 on zEC12

Proof Positive: zEC12 Performance Improvements

Predictable Linear Scalability for Operational BA

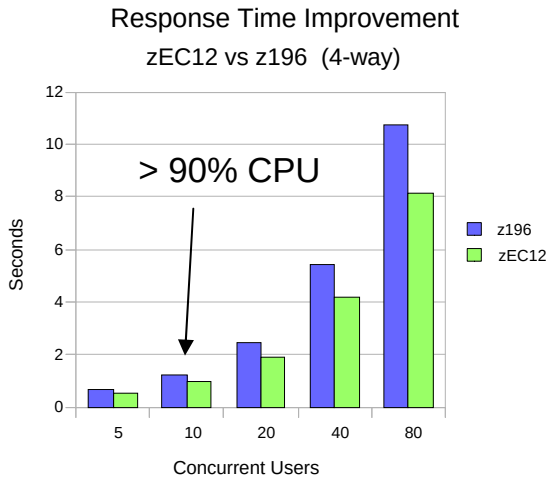

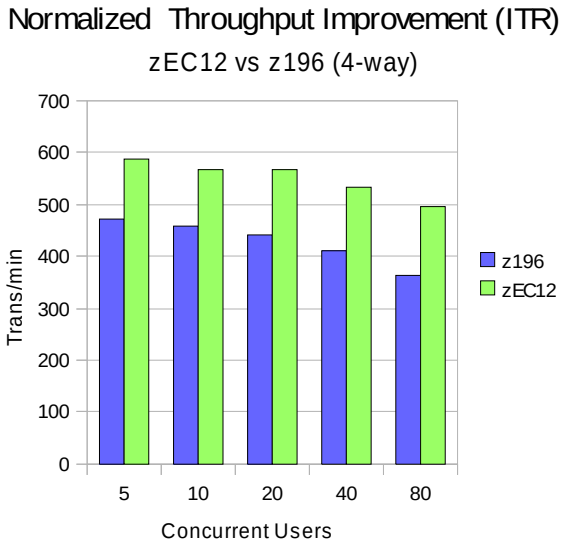

Scale to support more users in a linear predictable fashion

- 18% to 24% response time improvement with zEC12
- 23% to 36% throughput improvement with zEC12

Scale to support more operational users without an increase in hardware investment

✓ System z improvements continue to drive value for analytic workloads

Proof Positive: zEC12 Performance Improvements
Predictable Linear Scalability for Batch Reporting

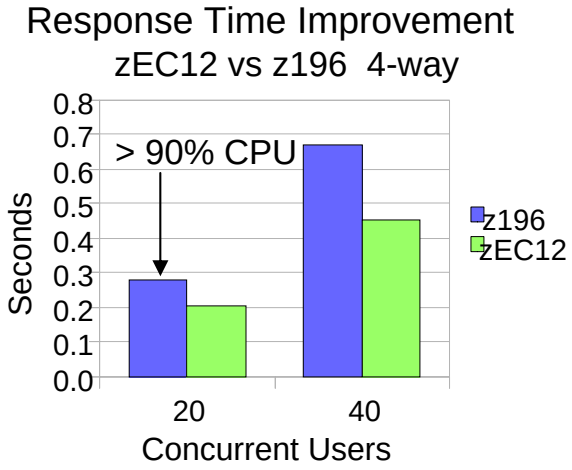

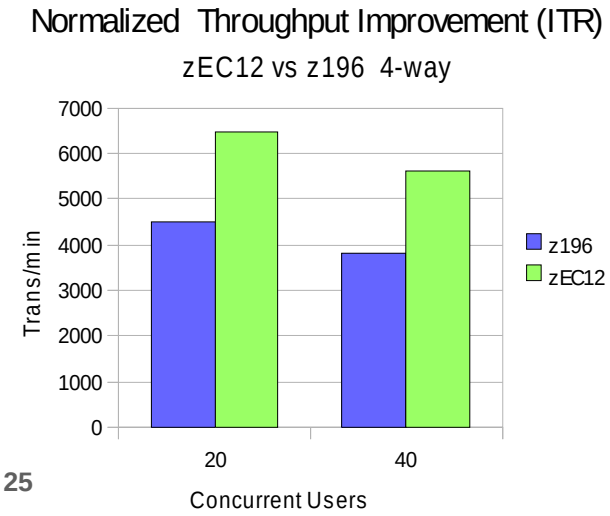

Scale to support more users in a linear predictable fashion

- 27% to 32% response time improvement with zEC12
- 43% to 48% throughput improvement with zEC12

Support your user community with sub second response time and improve your service level agreements.

✓System z improvements continue to drive value for analytic workloads

Cognos BI High Level Architecture

Cognos BI
High Level Architecture

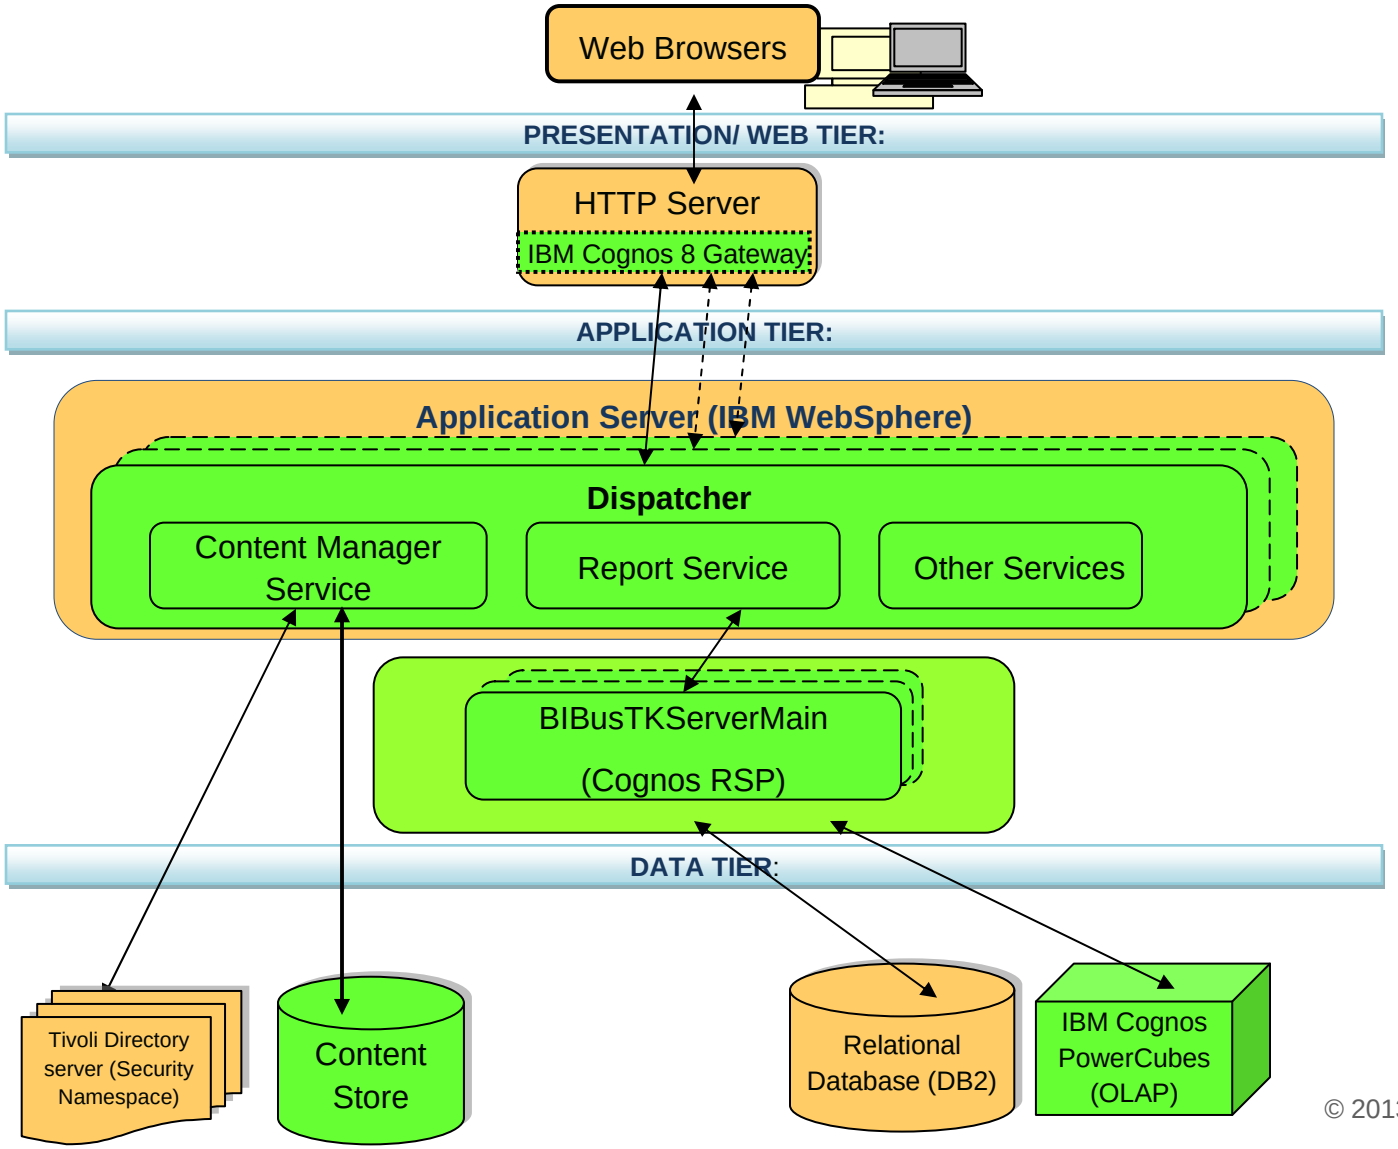

Cognos BI
High Level Architecture on z/OS –
Recommended

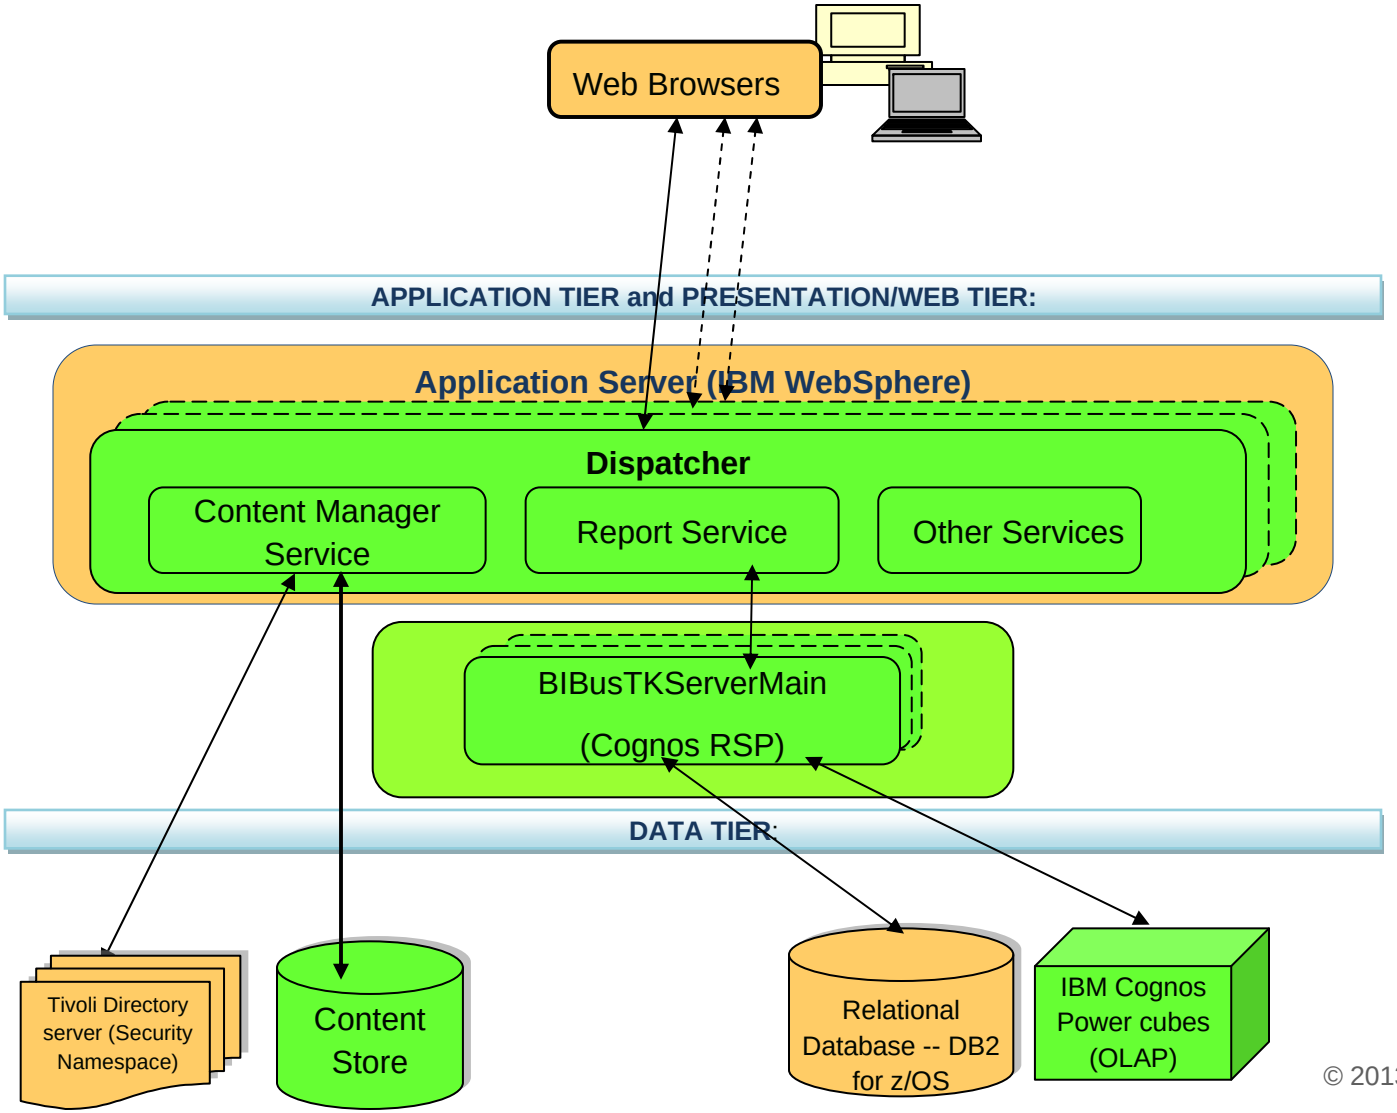

Cognos BI for zEnterprise

Install Steps

IBM zEnterprise® EC12 with IBM zEnterprise BladeCenter® Extension (zBX) Model 003

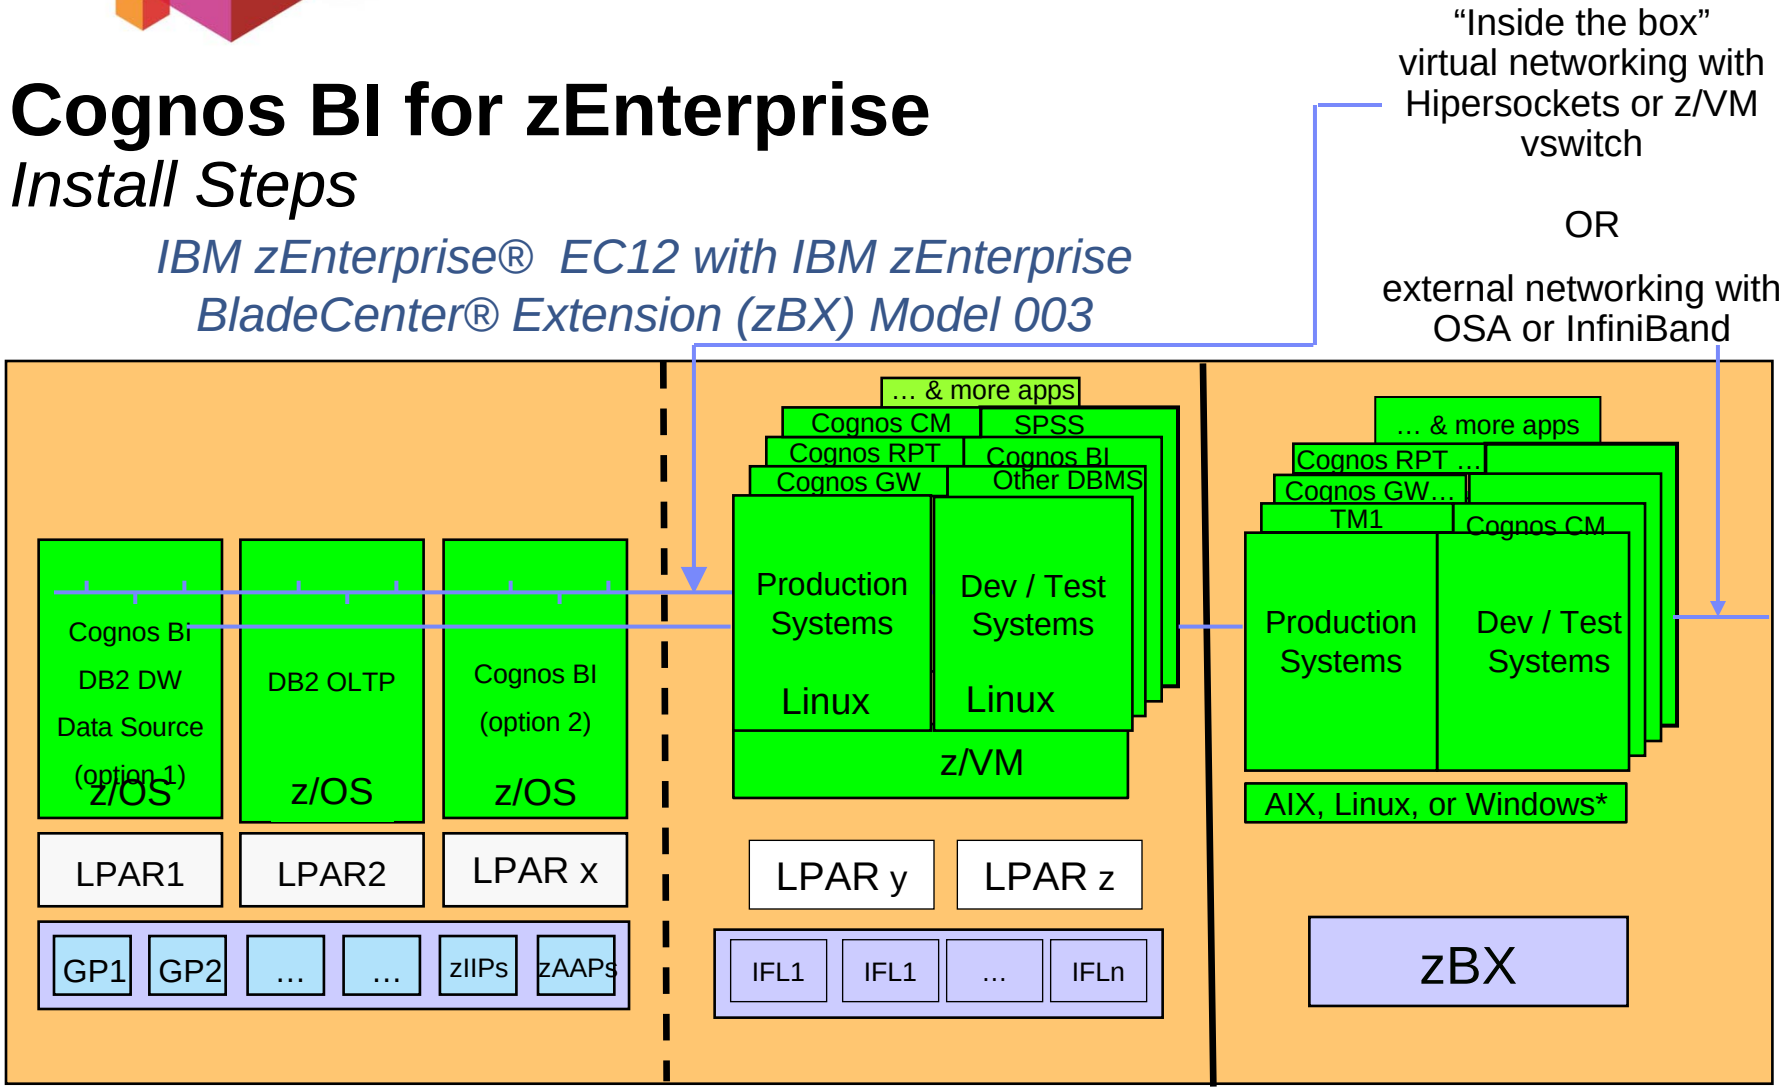

...a potential source of cost savings given zEnterprises ability to over-commit CPU capacity and automate with the Unified Resource Manager

IBM Cognos BI for zEnterprise Install Steps

…a potential source of cost savings given z/VM's ability to over commit CPU capacity

IBM zEnterprise® EC12 with IBM zEnterprise BladeCenter® Extension (zBX) Model 003

"Inside the box" virtual networking with Hipersockets or z/VM vswitch

OR

external networking with OSA or InfiniBand

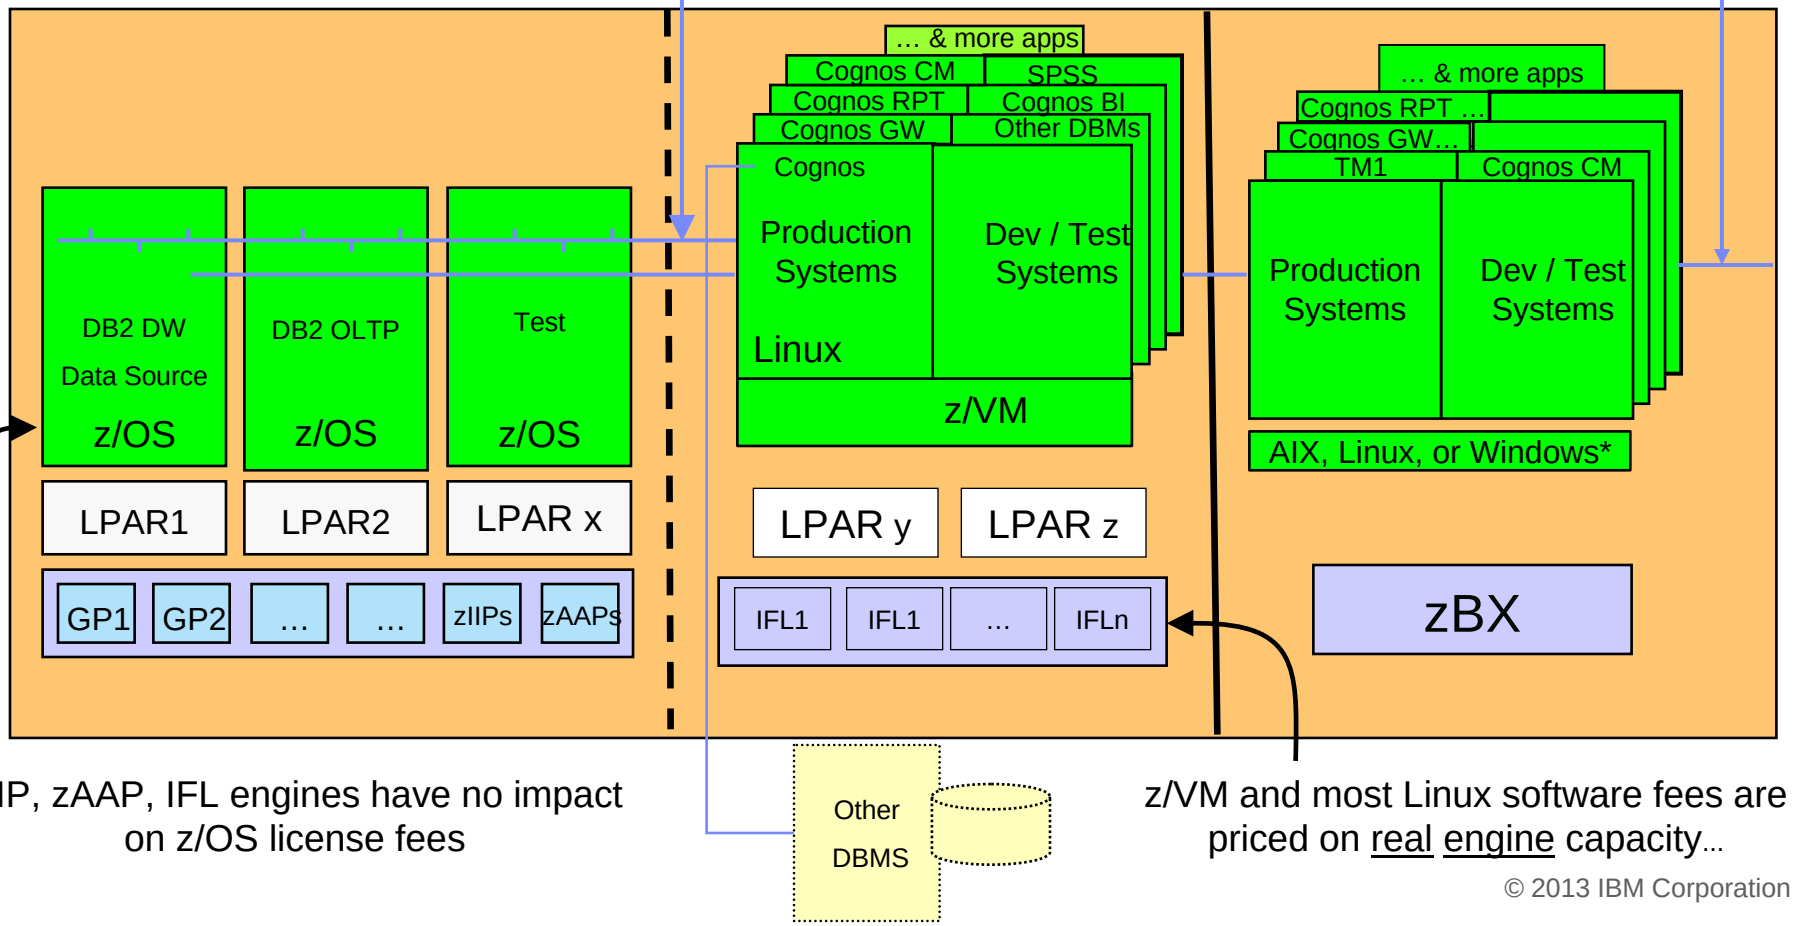

zIIP, zAAP, IFL engines have no impact on z/OS license fees

z/VM and most Linux software fees are priced on real engine capacity…

Cognos BI for zEnterprise

Implementing Cognos BI failover and managing growth option

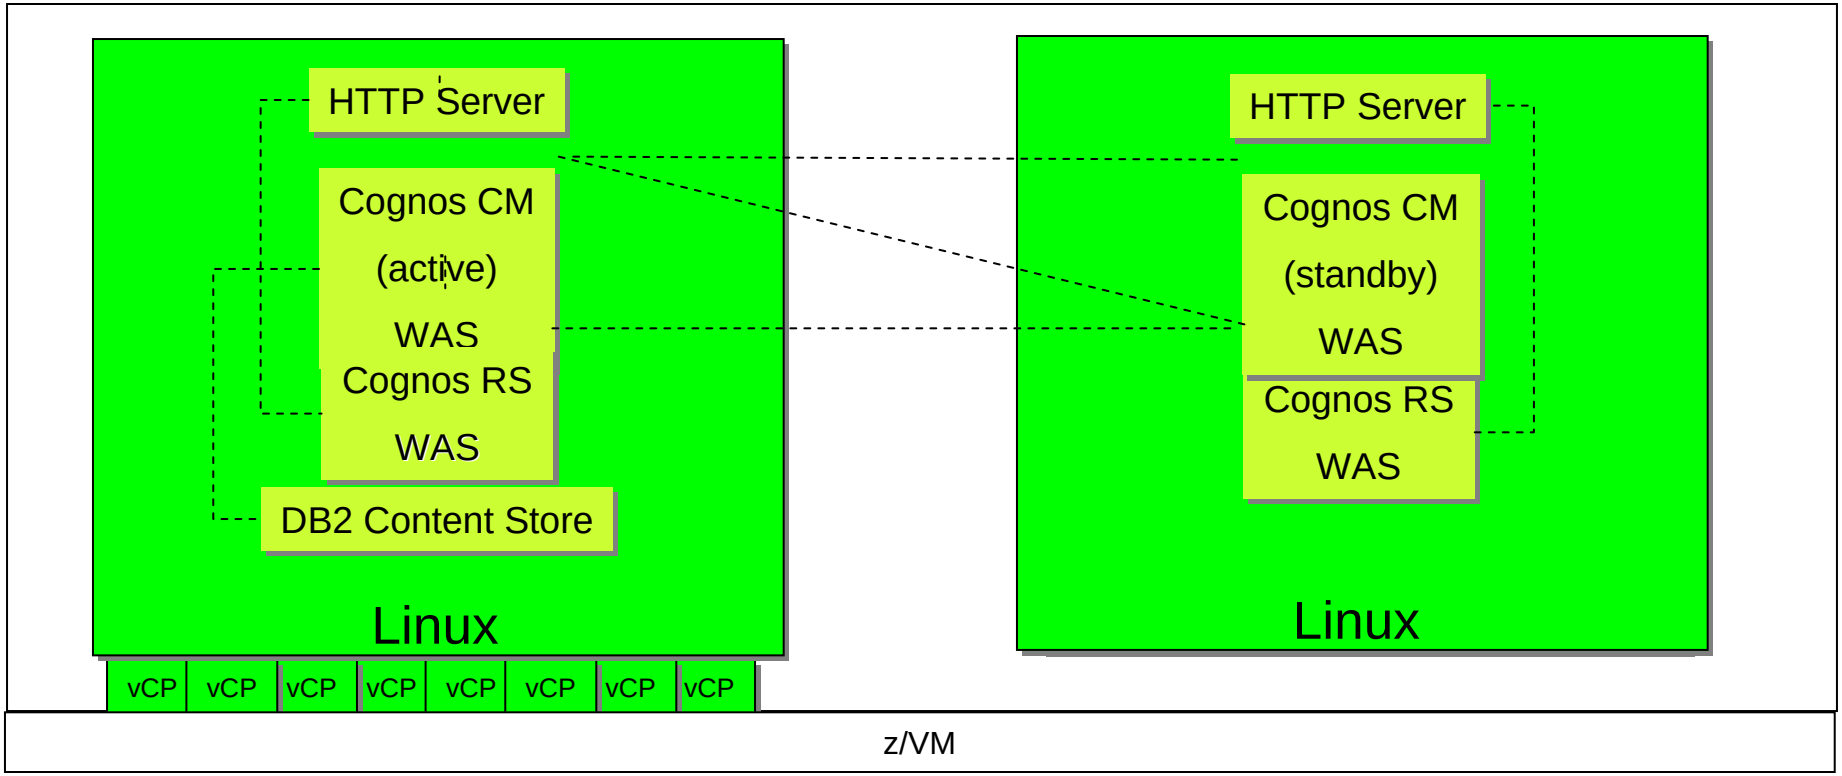

- **Vertical growth may be least cost – how large is “optimal”?**
- **What usage monitoring it needed?**
- **What SLAs have been committed?**
- **What components should be architected in failover?**

Predictive Analytics for Linux on zEnterprise

NEW

SPSS Decision Management for Linux on System z	SPSS Collaboration and Deployment Services for Linux on System z
▪ Employs both predictive models and business rules to automatically generate recommended actions	▪ Provides role-based models and security for in scoring, job scheduling, repository services, and integration
IBM SPSS Statistics for Linux on System z	**IBM SPSS Modeler for Linux on System z**
▪ Apply math to decision making and research for commercial, government, and academic users	▪ Data mining tool used for generating hypotheses and scoring ▪ Text analysis for unstructured data to model consumer behavior ▪ **In-Transaction Scoring with DB2 z/OS: Embeds the Scoring Algorithm Directly within the Transactional Application**

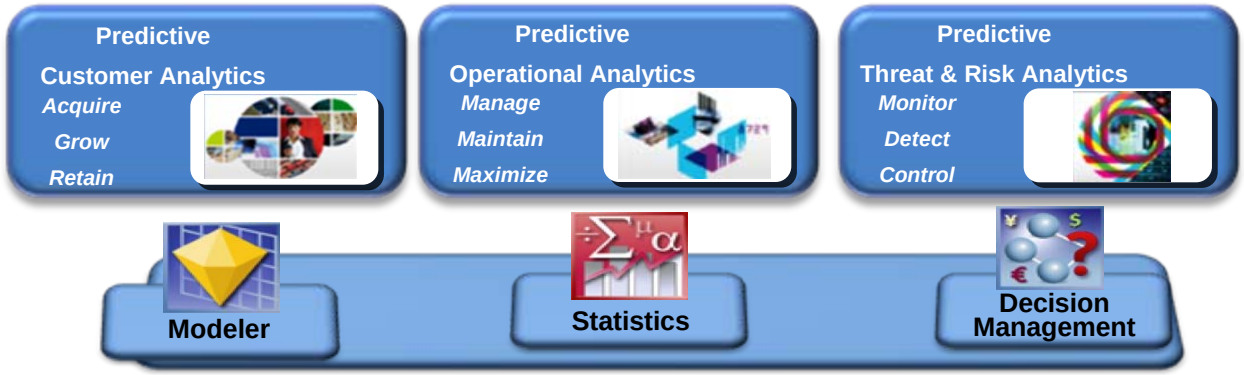

Scoring Options with an OLTP Application

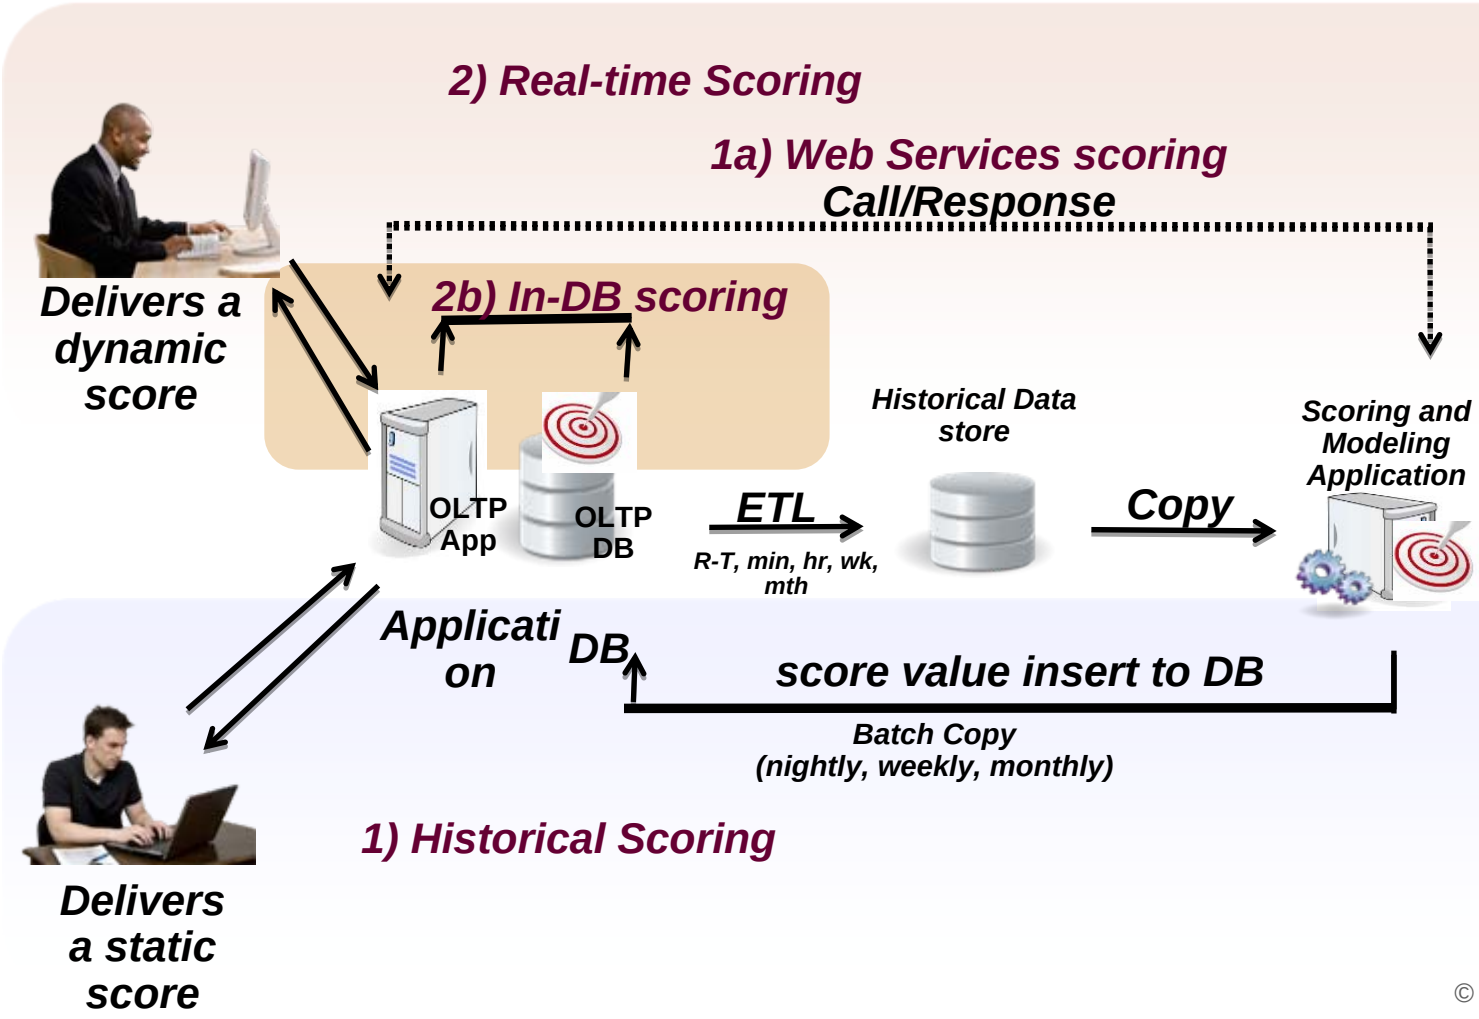

Modeler 15 Real-time Scoring with DB2 for z/OS

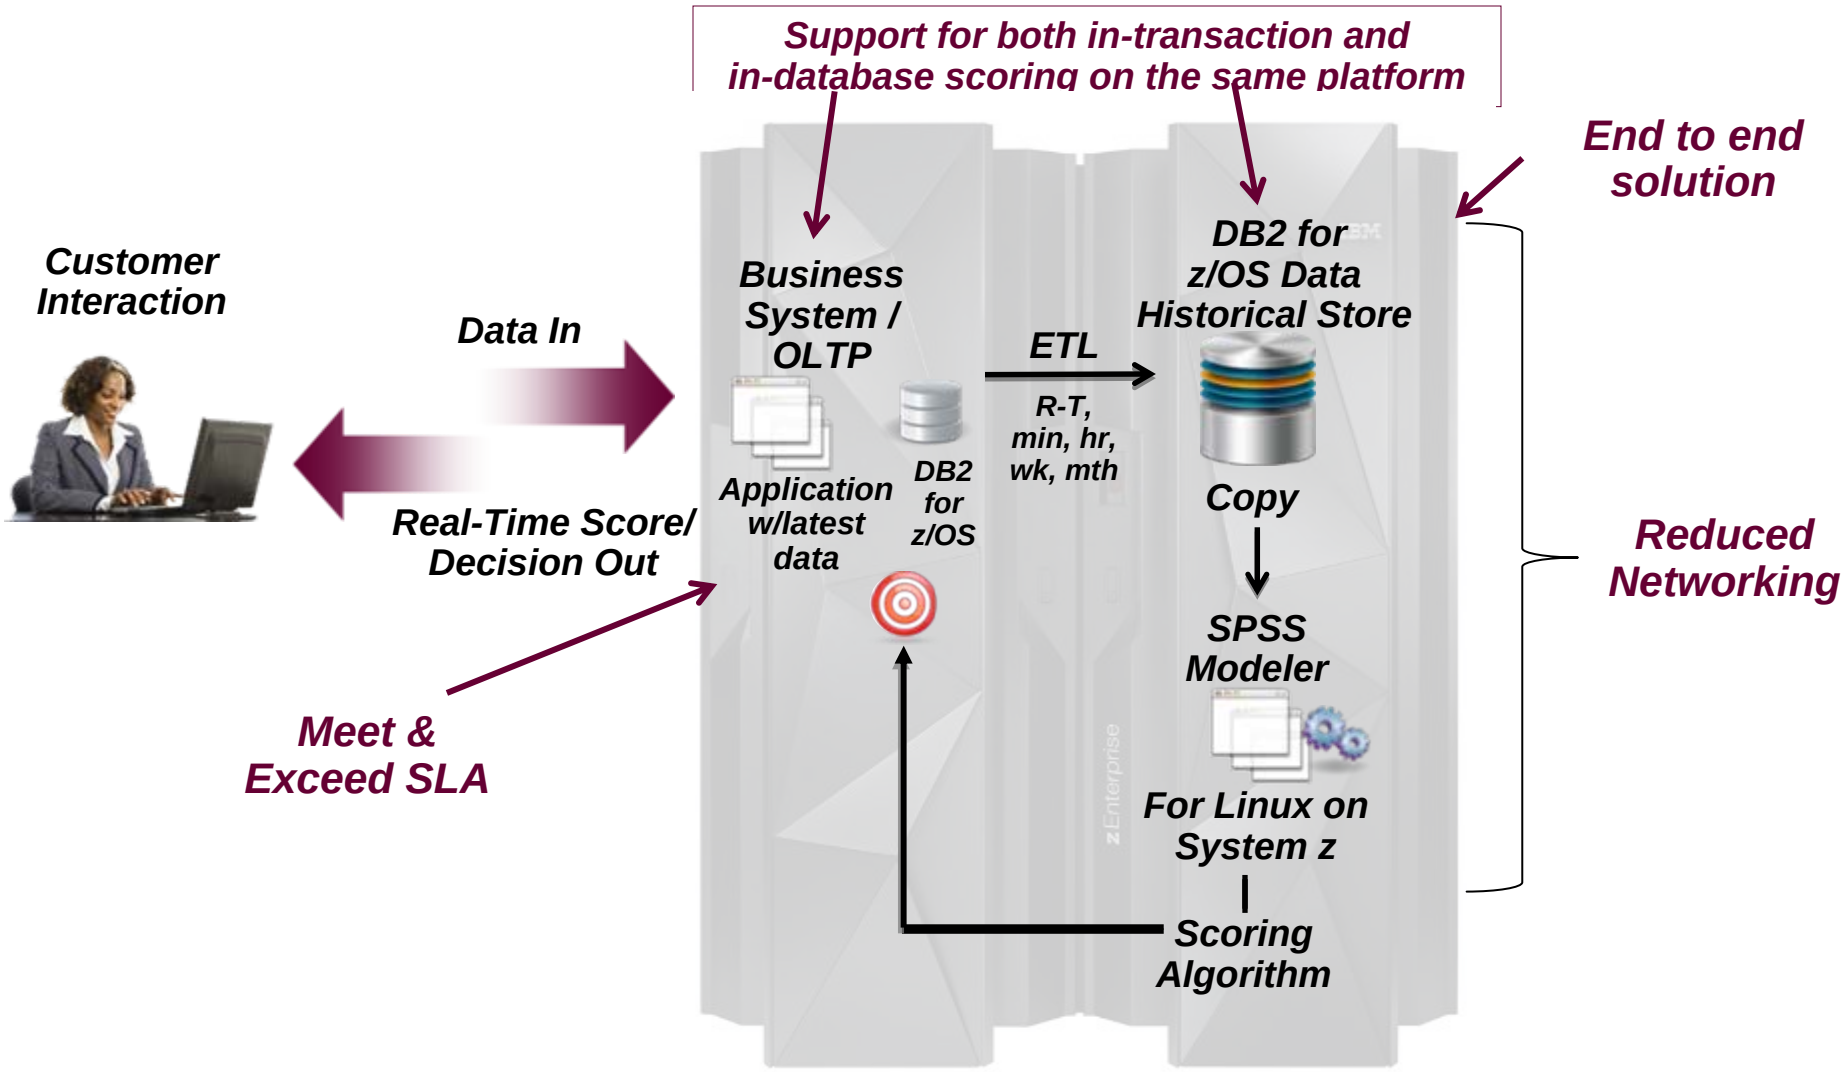

Can you really scale to support the volume of transactions an OLTP application processes without impacting performance?

- ✓ **Meets most demanding workload**
  - 3K to 5K transactions per second requested
- ✓ **Meets stringent SLA requirement**
  - Small additional CPU cost to score

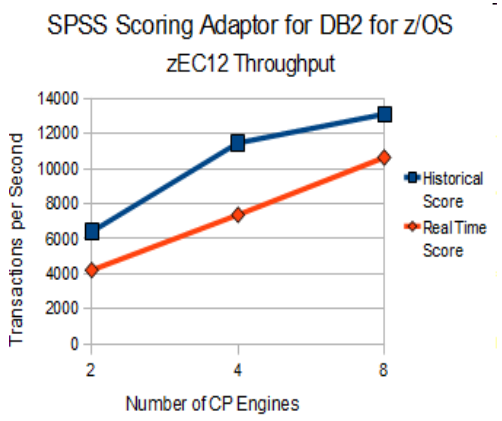

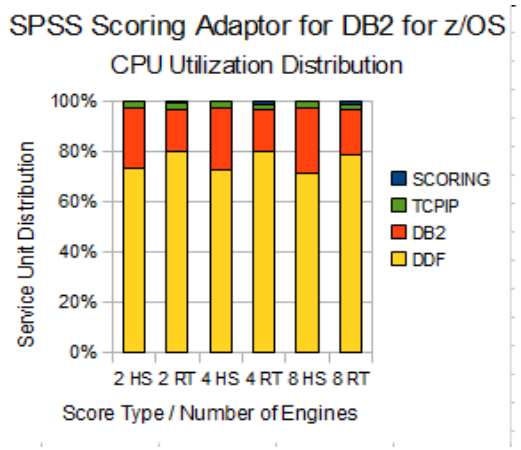

- ✓ **Provides best value**
  - Most accurate score is calculated in real time

SPSS Modeler Scoring Adaptor Real-Time Scoring z196 vs. zEC12 Performance Comparison

✓ **Meets most demanding workload**

- ~ 50% improvement in transactions per second
- at high utilization levels

✓ **Meets stringent SLA requirement**

- ~50% improvement in elapsed time
- at high utilization levels

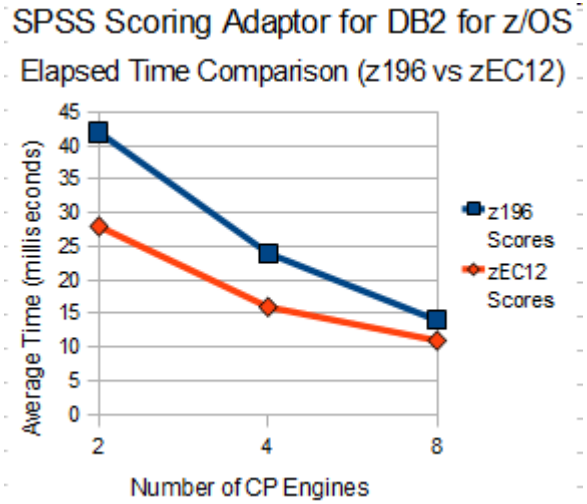

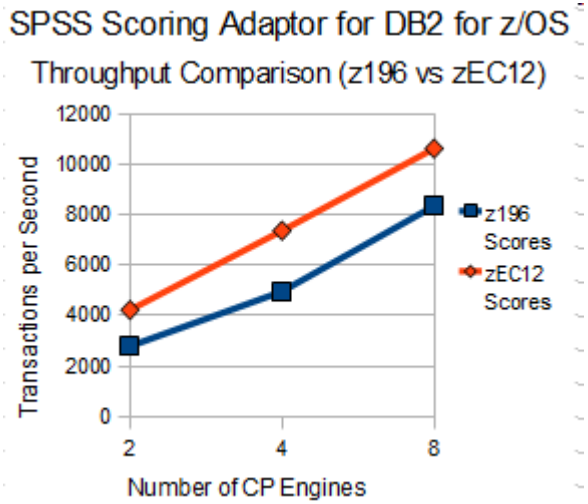

✓ System z improvements continue to drive value for analytic workloads

An enterprise information hub on a single, integrated, secure platform

Best OLTP/ Transactional Solutions

Industry leader in mission critical transactional systems

Best In Analytics

Industry recognized leader in BI, PA & DW solutions

Best In Flexibility

Start with your most critical business issue & quickly realize biz value with the flexibility to expand & grow across the enterprise

Transaction Processing (OLTP)

Business Analytics

Predictive Analytics

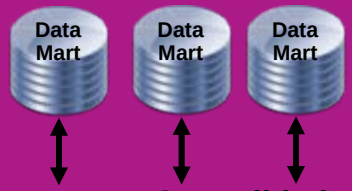

Data Mart Consolidation

zEnterprise

Recognized leader in workload management with proven security, availability, recoverability

DB2 Analytics Accelerator for z/OS *powered by Netezza technology*

Recognized leader in cost-effective high speed deep analytics

Unprecedented mixed workload flexibility & virtualization providing the most options for cost effective consolidation

IBM zEnterprise Analytics System 9700/9710

Mixed Workloads for Next Generation Business Analytics

The next generation of zEnterprise analytics; an integrated solution of hardware, software and services that enables customers to rapidly deploy cost effective game changing analytics across their business.

Preselected

All the necessary components are identified and integrated into an end-to-end solution

Pretested

Over 20 different customer-typical configurations are presized and tested

Solution Priced

Aggressively priced for a cost-effective add-on or new deployment for customers with critical data operations

Learn More!

- *Visit the IBM Mainframe Business Analytics & Data Warehousing Website*
  http://www.ibm.com/software/os/systemz/badw/

- *Join the IBM Analytics Networking Community*
  http://db2forzos.ning.com/group/datawarehousebusinessintelligenceonsystemz

IBM

THANK-YOU