Delivering Innovative Solutions to Meet Current and Future Market Demands

Ray Jones
Vice President, System z Worldwide Software
IBM Systems

New marketplace dynamics will drive hyper growth opportunity for IT Infrastructures

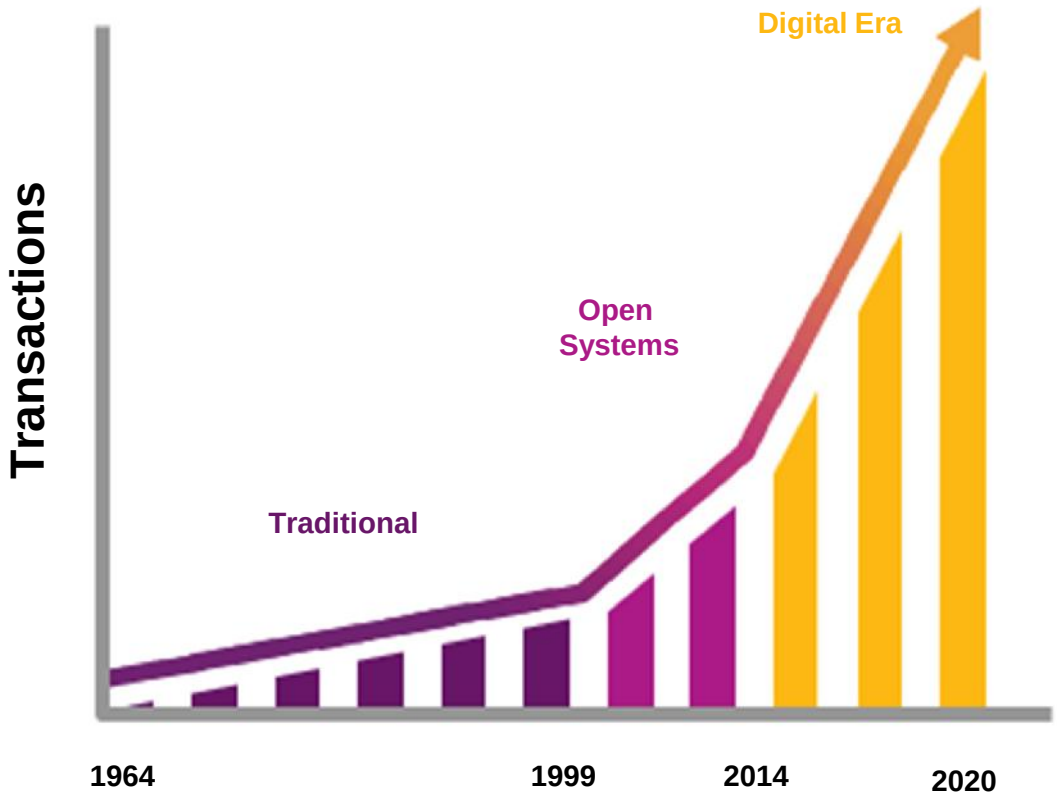

Traditional Workloads

1964–2014

- Batch
- General Ledger
- Transaction Systems
- Client Databases
- Accounts payable / receivable
- Inventory, CRM, ERP

Linux & Java

1999–2014

- Server Consolidation
- Oracle Consolidation
- Early Private Clouds
- Email
- Java, Web & eCommerce

Digital Era-Cloud Analytics Mobile Social

2015–2020

- On/Off Premise, Hybrid Cloud
- Big Data & Analytics
- Enterprise Mobile Apps
- Security solutions

Performance delivered through multiple dimensions

- **2X** performance boost for cryptographic coprocessors
- Symmetrical multi threading
- SIMD vector processing
- **2X** increase in channel speed
- **2X** increase in I/O bandwidth
- **3X** increase in memory
- **2X** increase in cache

- More scale for mobile transactions
- Faster data sharing between systems
- Real time predictive analytics
- Faster decision making with data-in-memory
- Lower Cloud cost

Java 8 and z13: Optimized CICS, IMS, and DB2 transactions

Up to 50%
improvement for generic applications

Up to 2X
improvement in throughput per core for security enabled applications

- **Up to 76% improvement** in throughput from z13 (SMT, SIMD, CPACF and more)
- **Up to 42% improvement** in throughput from IBM Java 8
- **Up to 60X improvement** with Java 8 exploiting z13 new SIMD vector hardware instructions for specific Java libraries and functions

Java 8

Compilers: z13 exploitation for increased performance

New

Enterprise COBOL for z/OS v5.2

- Leverage SIMD instructions to improve processing of certain COBOL statements
- Increased use of DFP instructions for Packed Decimal data
- Support COBOL 2002 language features: SORT and table SORT statements
- Allows applications to access new z/OS JSON services

Up to **14%** reduction in CPU time*

Enterprise PL/I for z/OS v4.5

- Leverage SIMD instructions to improve code for SEARCH and VERIFY
- Raised string size from 32K to 128M
- Improved middleware support
- Provide full support for JSON (Parse, Generate, and Validate)
- Addressed 28 RFE's

Up to **17%** reduction in CPU time*

z/OS XL C/C++ V2R1M1

- Vector/SIMD support (option, datatype, and built-in functions)
- High performance Math Libraries specifically tuned for z13
- New support for inline assembly (GNU compatibility)
- New Debug support for Vector/SIMD data type

Up to **17%** increase in throughput*

* The performance improvements are based on internal IBM lab measurements. All benchmarks were optimized and executed on zEC12 and z13, and built using the highest optimization level. Performance results for specific applications will vary, depending on the source code, the compiler options specified, and other factors.

CICS and z13

IBM z13: built for the needs of digital business and designed for the mobile generation

24% reduction in CPU per transaction seen in CICS on z13 for internal CICS VSAM benchmark compared with zEC12

31% improvement in ITR for the same workload

35% reduction in CPU/tran for an internal workload using SSL over HTTP with persistent sessions, leveraging z13 CPACF

Over 30% reduction in CPU/tran in CICS for internal CICS-DB2 benchmark

Please Note: Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

IBM z13 and DB2 z : Compete on Speed to Insight and Action

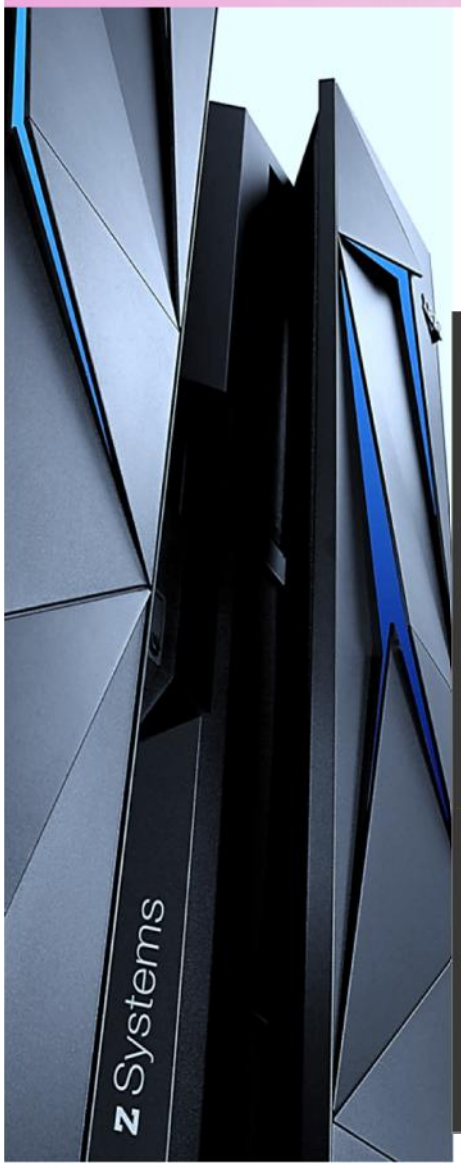

- DB2 workloads showing up to 38% range (mostly better than expected 10% improvement)
- 2.4x Reduction in compression cost in utility

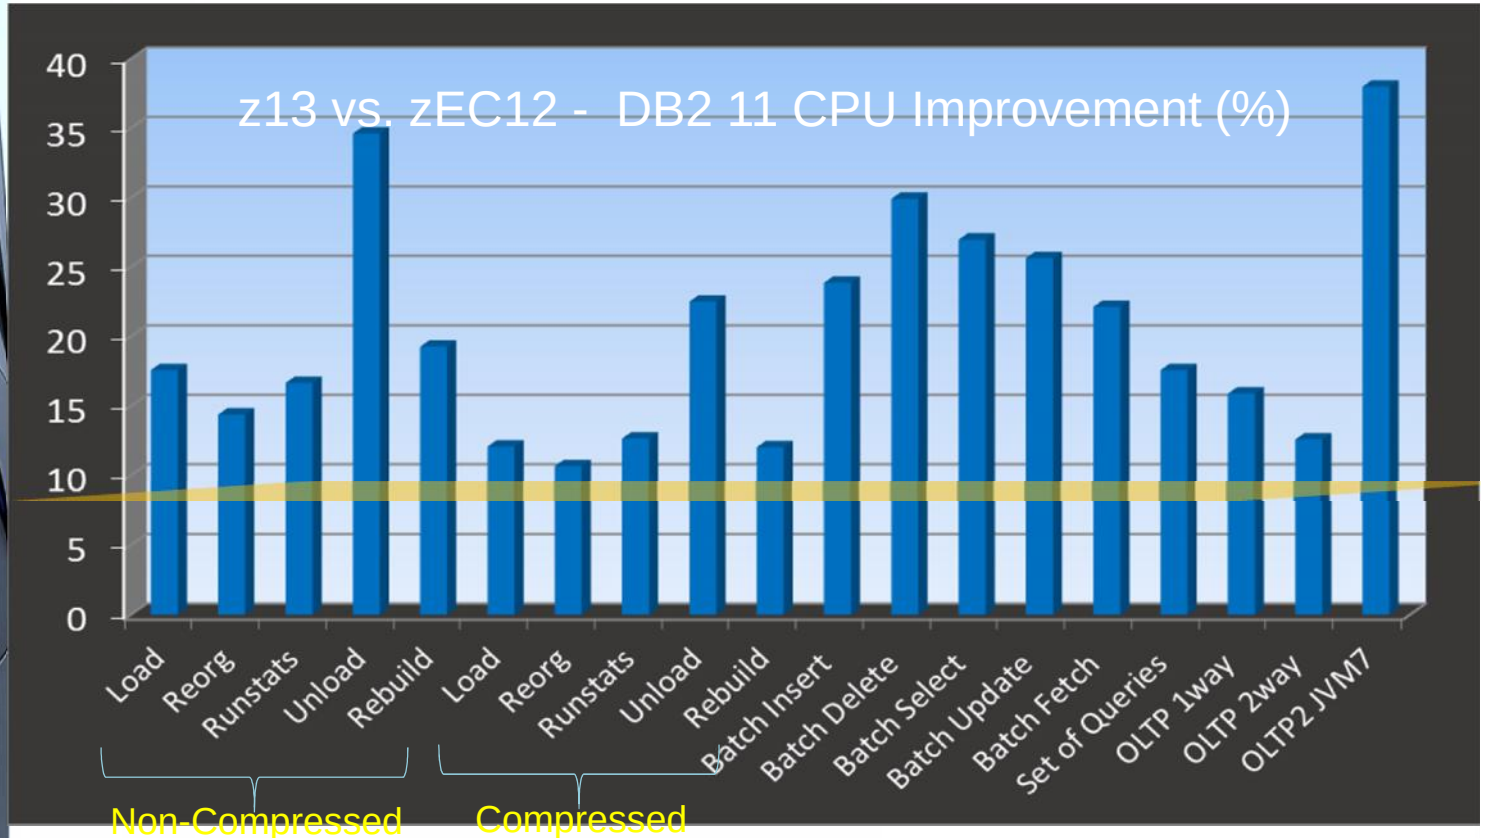

IBM z13 and DB2 z : Up to 4TB Memory

- DB2 workloads can take advantage of larger memory in z13
- Example : Banking transaction workload showing 24% throughout improvement by expanding DB2 buffer pools from 161GB to 638GB

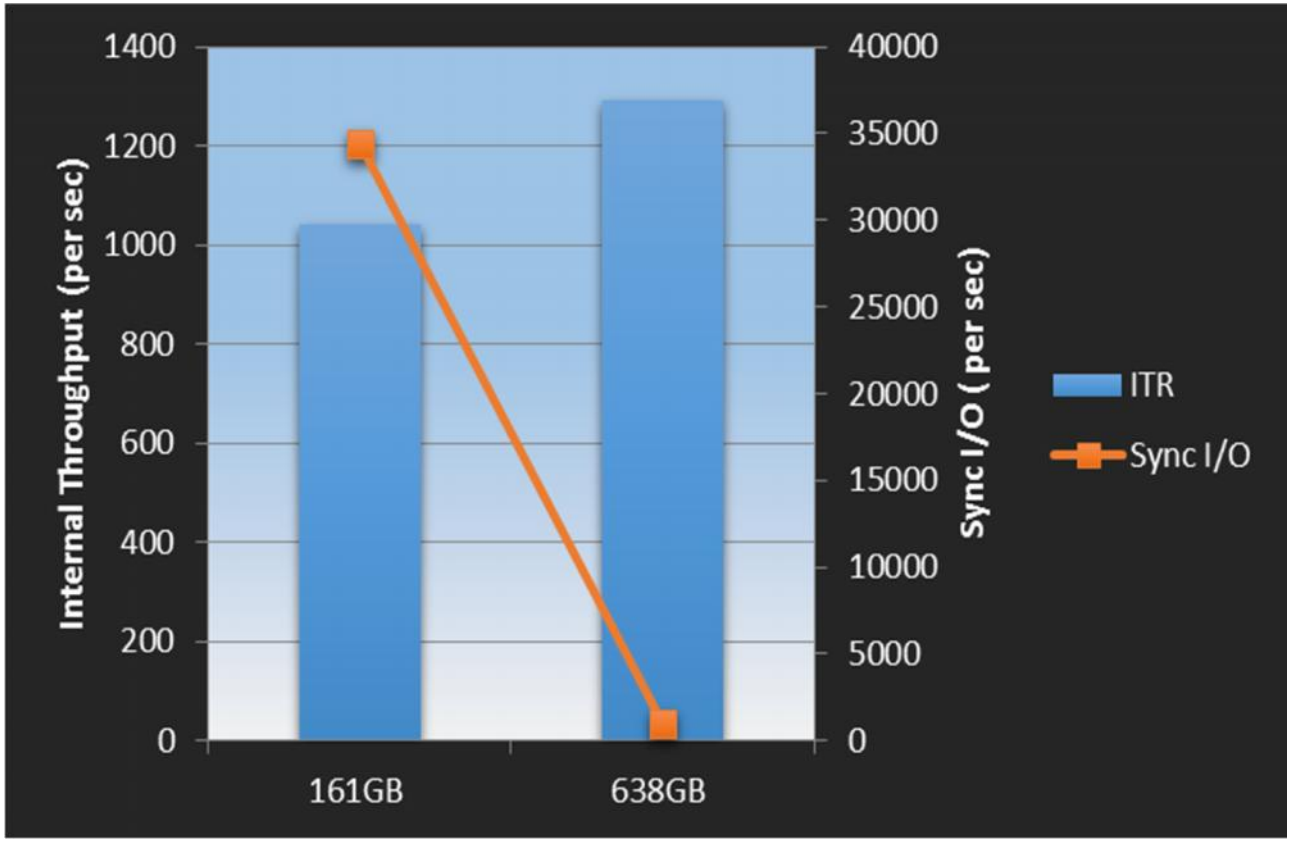

IMS 13 Fast Path Performance on z13

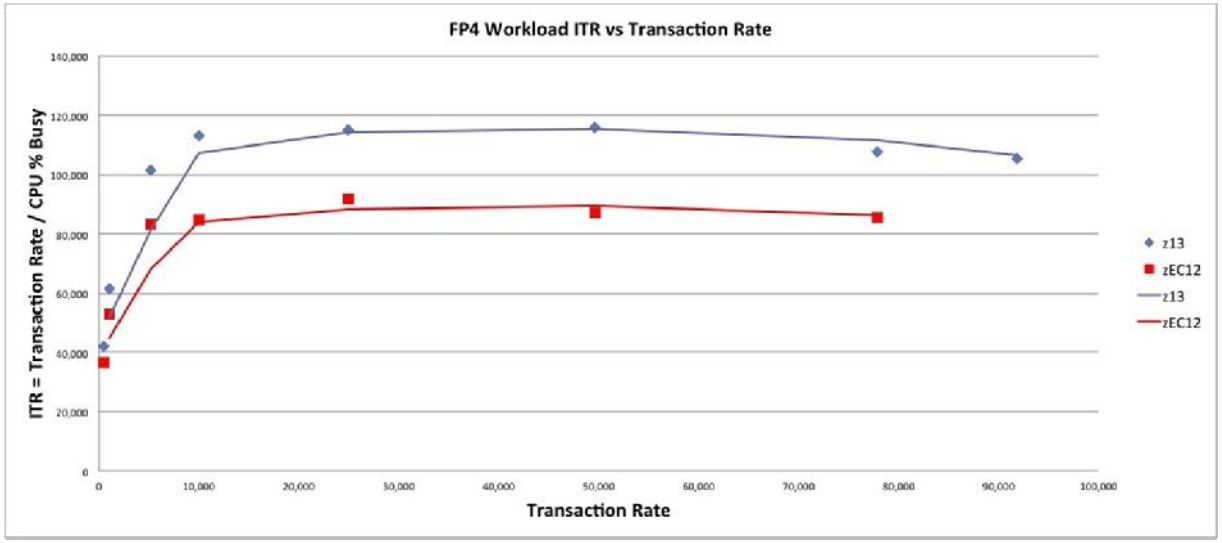

15-34% throughput improvement

over IMS 13 on equivalent EC12 using an IBM internal high volume fast path workload

•A 15-34% improvement in ITR was observed when comparing an IBM internal high volume IMS Fast Path workload running on a zEC12 vs z13 both configured with 9 general purpose engines

IMS 13, when running the IMS Java transaction processing workload on IBM z13, showed as much as a 38% increase in throughput at equivalent CPU as compared to zEC12

z/OS V2.2 and IBM z13 Provide Synergies

Meet the challenges of cloud, analytics and mobile workloads

The world's premier transaction and data engine enabled for the mobile generation

Driven by security, resiliency, and the economics of scale

- Strengthened security – encrypt **2X** as fast with CPACF on IBM z13™ (z13)*
- Signed audit records to help improve compliance
- Crypto cards now shareable by 85 LPARS – more than a **5X** increase

An integrated transaction and analytics system for real-time insights

Powered by data serving, analytics, powerful batch

- Selected key z/OS ATLAS 3.10.0 functions are accelerated using SIMD instructions and demonstrate up to **80%** higher throughput on z13 than on zEC12.***
- For eligible data, store up to **4X*** more data with zEDC
- Improve performance of many mobile, cloud, and analytics applications running on z13 with SIMD with Java® SDK 8

The world's most efficient and trusted cloud system that transforms the economics of IT

Fueled by server scale, large memory, high availability and resiliency

- **141-way** support on z13, approximately **40%** more cores than IBM zEnterprise® EC12 (zEC12)
- Up to **4TB** memory
- An average capacity improvement of **38%** compared to zEC12 Including use of SMT for zIIPs**

- **Improved Simplification**
- **z/OSMF now included with z/OS**
- **Easier operations**
- **Superior economics for an improved migration period**

*These results are based on projections and measurements completed in a controlled environment Results may vary by customer based on individual workload, configuration and software levels

The z13 provides lower overall mainframe costs through the ability to process more workload on larger zIIPs with an average capacity improvement of **38% compared to zEC12 including the exploitation of the new multithread option on the z13 zIIP.

***This claim is based on results from internal lab measurements. The double precision function improvement is derived from comparisons of a select set of commonly used z/OS ATLAS 3.10.0 functions executing on z13 to the equivalent functions executing on zEC12. A subset of these functions is accelerated using SIMD instructions on z13 The SIMD benefit is demonstrated using this subset. The performance improvements achieved will vary depending on the workload and other factors. .

zLinux – More Scale, Availability and Open Source

New

SMT technology on z13	Up to **38%** improvement in core performance	Up to **20%** more virtual servers per core
IBM Spectrum Scale	Disaster recovery solution for mission-critical workloads	
KVM	New industry-standard hypervisor (SOD)	
zAware for Linux on z	IT Analytics for improved availability	
Elastic Storage for Linux on z Systems	Enables new class of workloads	
PostgreSQL and Docker	Open support extended with PostgreSQL and Docker support	

Mobile workloads Impact Systems more than web workloads.

- Increase in peak and off-peak transactions. Expect 10-50% growth as you add a mobile channel.
- Increased query or “read-only” transactions. As many as 50% of mobile transaction could be “read-only”.
- Unanticipated spikes in workload due to popular apps, features or special offers. No traditional times for workload spikes.
- Inefficient applications written by “non-professional” coders. Drives up transaction rates.

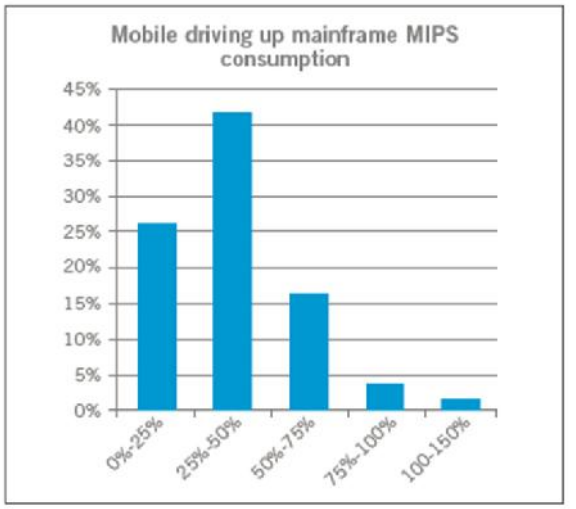

Chart 7: The increased use of mobile applications has increased MIPS consumption by more than two-fifths (41 percent), with 2 percent saying it has more than doubled.

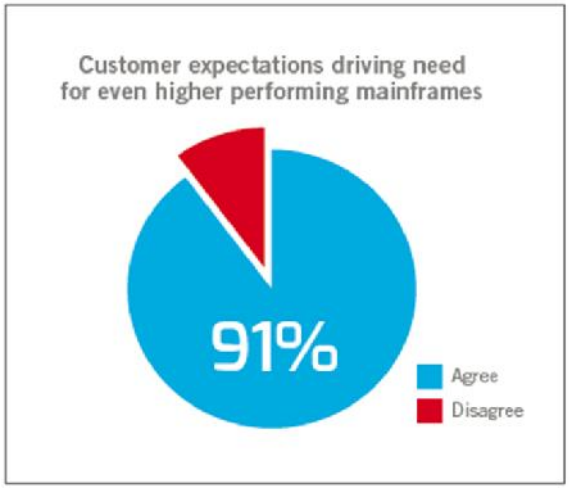

Chart 4: 91 percent of CIOs say now that customer-facing applications are using the mainframe and performance expectations on it have increased.

Solutions to accelerate connecting mobile devices to z Systems

- New — End-to-end reference architecture for mobile on z Systems from MobileFirst platform to CICS and IMS
- New — z Systems end-to-end mobile Security architecture
- New — IBM Mobile Center of Competencies worldwide (Poughkeepsie, Tokyo, Bejing, Boeblingen, Montpellier)
- New — Apple iOS apps package with GBS implementation services for z Systems
- New — Sample MobileFirst applications for CICS, IMS, z/OS Connect

Replicating z Systems Mobile Workloads increases TCA by 66% versus co-locating MobileFirst Platform and using Mobile Workload Pricing

Which platform provides the lowest TCA over 3 years?

Mobile Users

- **500 concurrent connections**
- **70% do 1 read/session; 25% do 4 reads/session; 5% do 20 reads/session with 100ms think time**
- **1 second cache invalidation**

Mobile read-only workload driving minimum throughput of **6,300** transaction per second and response time of 12ms

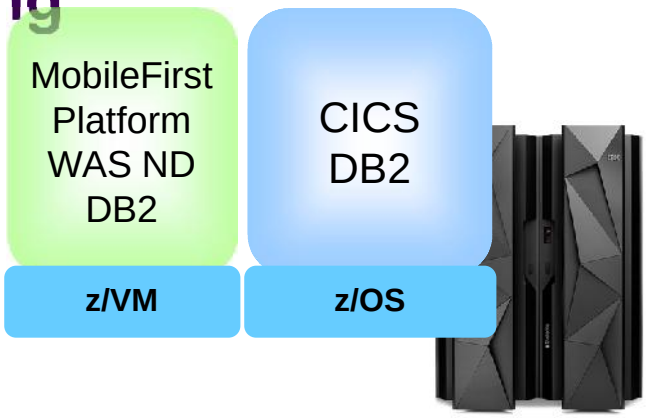

z13-606 + 4 IFLs*

$11.2M (3 yr. TCA)
Prod + Dev/QA + DR
Mobile Workload Pricing

Estimated **40%** lower cost for systems compared

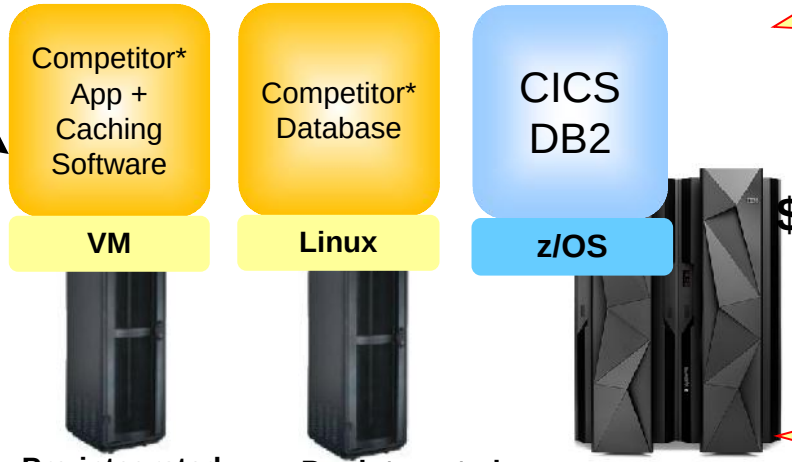

Pre-integrated Competitor* Eighth Unit (30 cores pro-rated)

Pre-integrated Competitor* Eighth Unit (24 cores)

z13-606

$18.6M (3 yr. TCA) Prod + Dev/QA + DR

Estimated **66%** higher cost for systems compared

* Competitor Caching and Database sizing estimated from WebSphere Extreme Scale Caching Test.
** Estimated performance, sizing and cost for z13 based on tests conducted on zEC12

This is based on an IBM internal study designed to replicate a typical IBM customer workload usage in the marketplace. Test involved executing a materially identical mobile transaction processing workload in a controlled laboratory environment with comparable tuning and sizing. Prices, where applicable, are based on US prices as of 12/31/2014 for both IBM and competitor. Price comparison based on 3 Year Total Cost of Acquisition (TCA) includes all HW, SW and 3 years of service & support. Sizing shown is for Production to which 30% is added for System z for Dev/QA and CBU pricing for DR and 2x for Distributed.

MobileFirst Platform on Linux on z System* is expected to provide lower front-end cost and better scalability than x86

At 50 concurrent users, z Systems provides better 3-year TCA

	MobileFirst Platform on Linux on z Systems*		MobileFirst Platform on x86	
# Concurrent Users	Front-end Cost per TPS	Response Time (ms)	Front-end Cost per TPS	Response time (ms)
10	$2,634	42	$2,074	50
30	$1,091	43	$1,066	54
50	$812	44	$964	62
100	$525	48	$770	68
200	$456	70	$636	95
400	$439	131	$693	205

16% better

37% better

Green = Better

* Estimated performance, sizing and cost for z13 based on tests conducted on zEC12

This is based on an IBM internal study designed to replicate a typical IBM customer workload usage in the marketplace. Test involved measuring throughput in transactions per second and response time for executing a materially identical mobile transaction processing workload in a controlled laboratory environment with comparable tuning and sizing. Prices, where applicable, are based on US prices as of 12/31/2014 for both IBM and competitor. Price comparison based on 3 Year Total Cost of Acquisition (TCA) includes all HW, SW and 3 years of service & support. Sizing shown is for Production to which 30% is added for System z for Dev/QA and CBU pricing for DR and 2x for Distributed.

Traditional Analytics Approach

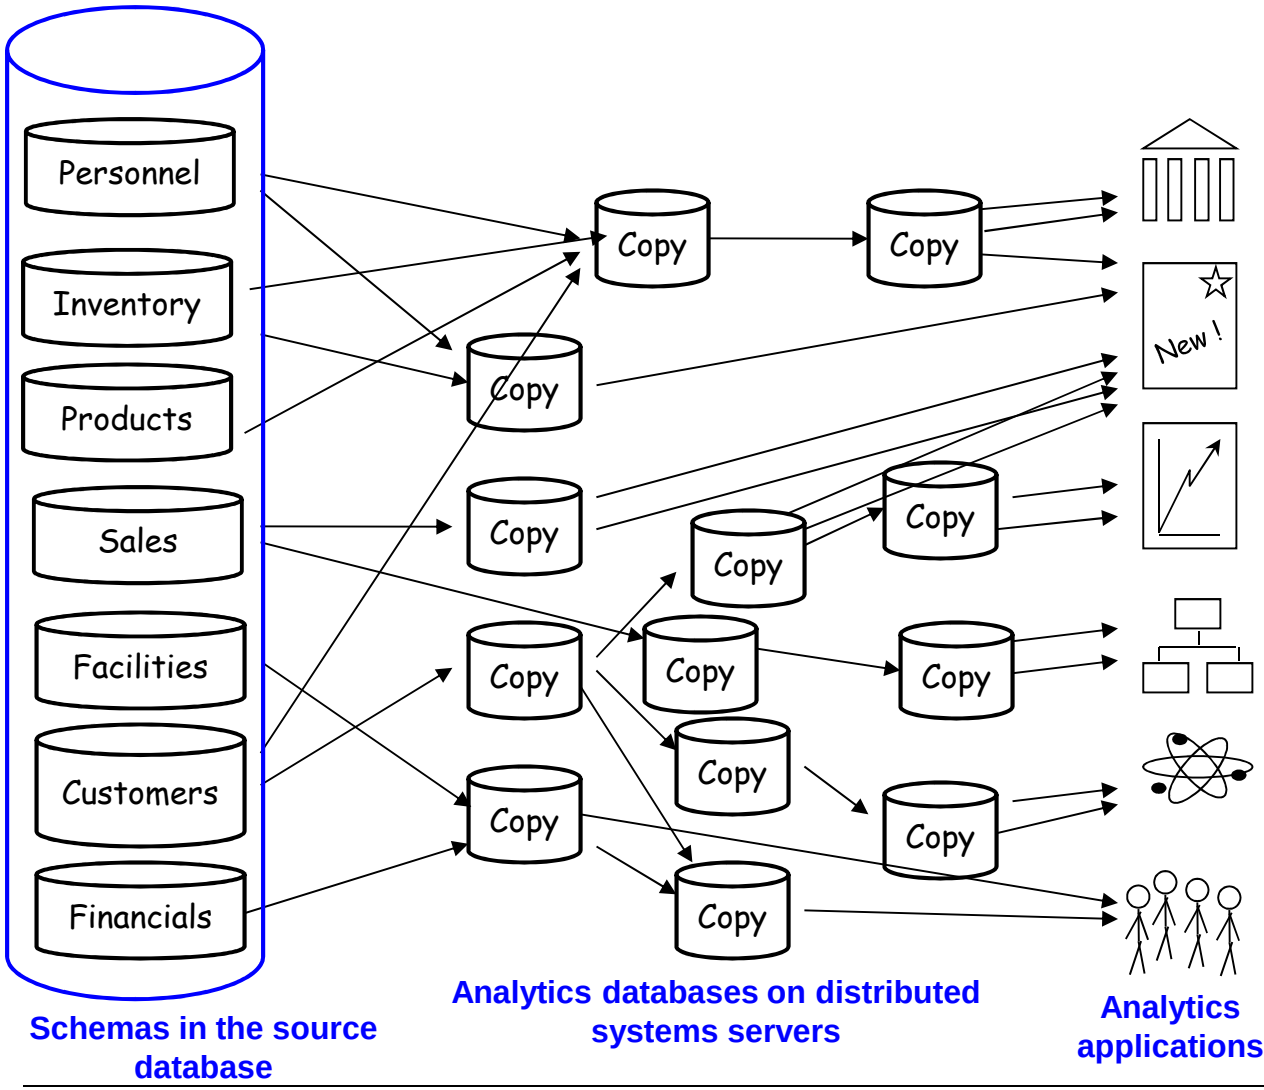

Problems:

- Data latency: time between transaction and insight
- Expensive, resource-intensive data replication processes
- Greater risk of data security breaches
- Data governance issues: copies of data can become inconsistent – do users trust the data?
- Data currency challenges: copies of data can become out-of-date – users demand timely data
- Proliferation of data silos impedes integration, reduces value derived from data assets

New applications and Data Sources
Health Scenario

Leverages:
Fast IFLs
Large shared memory
Data on z
Cloud deployment approach
More models tested
More iterations tested
More data included
Faster time to winning results

Other projects underway:
Gene Essentiality
Alzheimer's Disease
Heart Disease
Asthma
Prevent Falling
Breast Cancer

Rheumatoid Arthritis Responder DREAM Challenge

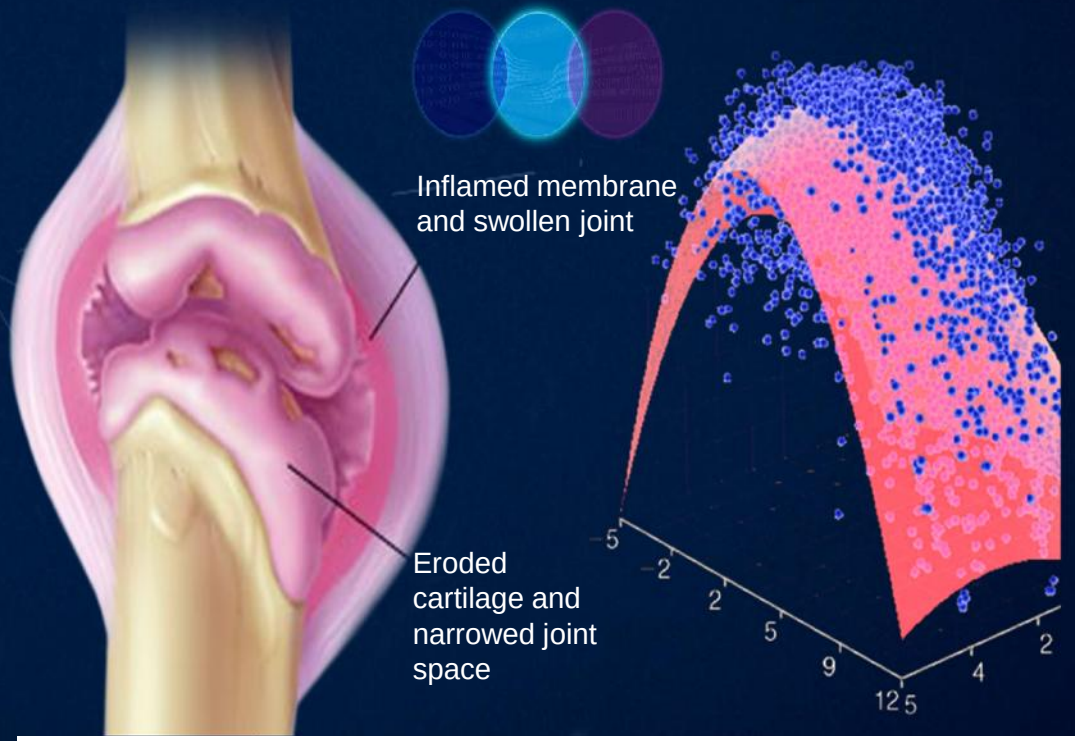

Sage · DREAM · CORRONA · ARTHRITIS FOUNDATION Take Control. We Can Help. · MERCK Be well · novo nordisk

Keeping the data on z13 and making a copy for DB2 Analytics Accelerator saves over 88%

Estimated 4 yr. cost summary

System costs = $1,052,901
Labor costs = $137,613
Total = $1,190,513

Assuming 4 cores on z13 running at 85% utilization and 140 x86 cores on N3001-010 running at 45% utilization, transfer will burn **260 MIPS** and use **0.44 x86 core** *per day*

This is based on an IBM internal study designed to replicate a typical IBM customer workload usage in the marketplace. Test involved measuring in a controlled laboratory environment elapsed time for system and administrator to extract, send and receive 1,118GB file from z13 to DB2 Analytics Accelerator N3001-010 (Mako Full Rack. Prices, where applicable, are based on US prices as of 12/31/2014 for both IBM and competitor. Estimated amortized cost from 4 Year Total Cost of Acquisition (TCA) that includes all HW, SW (OS, DB and tools) and 4 years of service & support. For Labor costs, used annual burdened rate of $159,600 for IT Administrator for z Systems and x86. Results may not be typical and will vary based on actual workload, configuration, applications, queries and other variables in a production environment. Users of this document should verify the applicable data for their specific environment.

z Systems Is Optimized For Operational Analytics

***17x** performance*
***13x** price performance!*
for systems compared

Standalone Pre-integrated Competitor V4

Eighth Unit

Workload Time	1,810 mins
Reports per Hour (RpH)	5,343
Competitor Eighth Unit (HW+SW+Storage)	$2,746,041

$514
Per Report per Hour (3yr TCA at no discount)

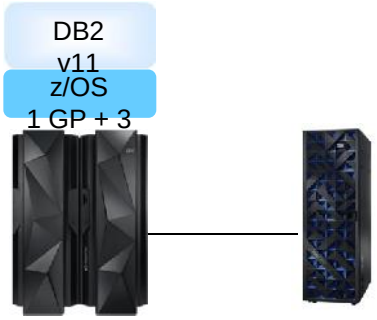

z13 + IBM DB2 Analytics Accelerator (N3001-010)

Workload Time	105 mins
Reports per Hour (RpH)	92,095
z13 (1 GP + 3 zIIP, HW+SW+ Storage) + Accelerator V4.1 with PDA N3001-010 hardware	$3,652,131

IBM zEnterprise Analytics System 9700

$40
Per Report per Hour (3yr TCA at no discount)

Based on IBM sponsored and internal tests comparing IBM zEnterprise Analytics System 9700 with a comparably priced, comparably tuned competitor Eighth Unit configuration (version available as of 12/31/2014), executing a materially identical 10 TB BIDAY "Fixed Execution" workload in a controlled laboratory environment. Test conducted with BIDAY "Fixed Execution" workload measures elapsed time for executing 161,166 concurrent reports using 80 concurrent users. Intermediate and complex reports are automatically redirected to IBM DB2 Analytics Accelerator for z/OS (powered by N3001-010 hardware or Mako). Price comparison based on a 3YR Total Cost of Acquisition (TCA) using U.S. prices current as of December 31, 2014, including hardware, software, and maintenance. Compared prices exclude applicable taxes, and are subject to change without notice. Competitor configuration: Eighth Unit including competitor recommended software options and features. IBM configuration: z13 platform with 1CP and 3 zIIPs with 128GB memory and DB2 Analytics Accelerator Full Rack (N3001-10) with 7 S-blades (140 Intel E5-2680v2 2.8GHz cores and 128 GB RAM), 2 Hosts (1 active – 1 passive) with 20 Intel E5-4650v2 2.4GHz cores each and 12 disk enclosures, each with 24 600GB SAS drives . Results may not be typical and will vary based on actual workload, configuration, applications, queries and other variables in a production environment. Users of this document should verify the applicable data for their specific environment.

In-Transaction: What About System Impacts?

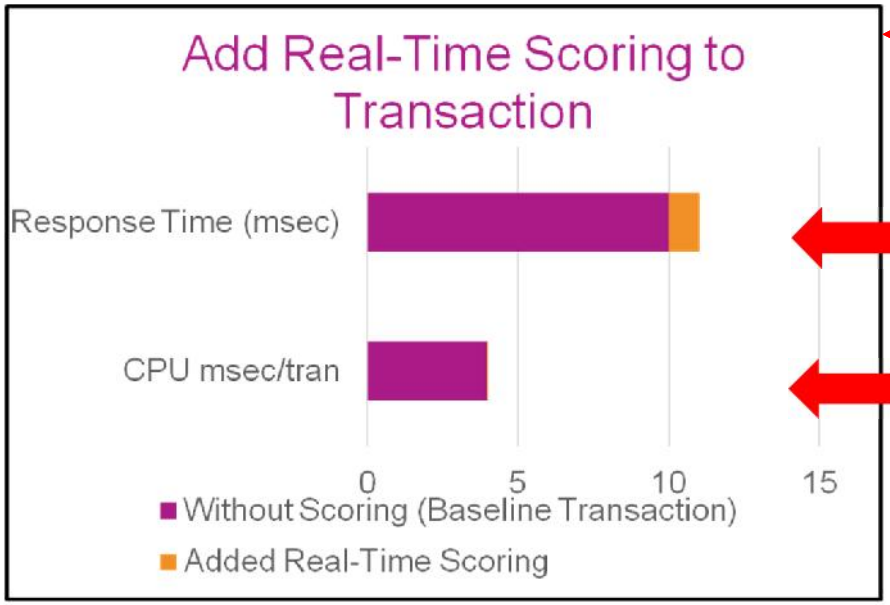

< 1 msec difference

0.06 msec difference

Minimal SLA and CPU impact

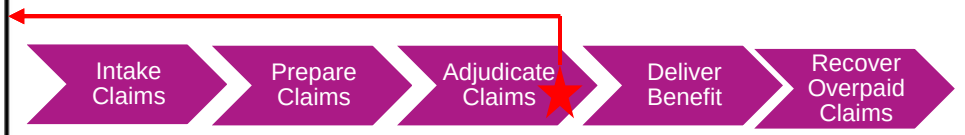

***Sample* Workload Specifics (Lab Measurements):**

IBM z13
Transactions > 320,000 in 5 min.
Predictive Model: Logistic Regression
Inputs to Model : 12

***** Testing done with IBM SPSS & DB2 z/OS***
***** Results will vary based on: modeling algorithm, number of inputs, amount, scope and type of data preparations***

***** Internal testing shows scoring to be a fairly fixed consistent cost, so relative impact will improve with transactions heavier than 10ms baseline***

This measure was done with a baseline business transaction lasting 10ms, but the relative performance cost of scoring will be even less with transactions that are longer, for example:

If a credit card authorization process takes 300ms, then the impact of in-transaction scoring would be < 1ms compared with the 300 ms transaction

If a claims transaction takes 2 seconds in transactional batch, then impact of in-transaction scoring would be < 1ms compared with the 2 seconds

What does in-process scoring mean?

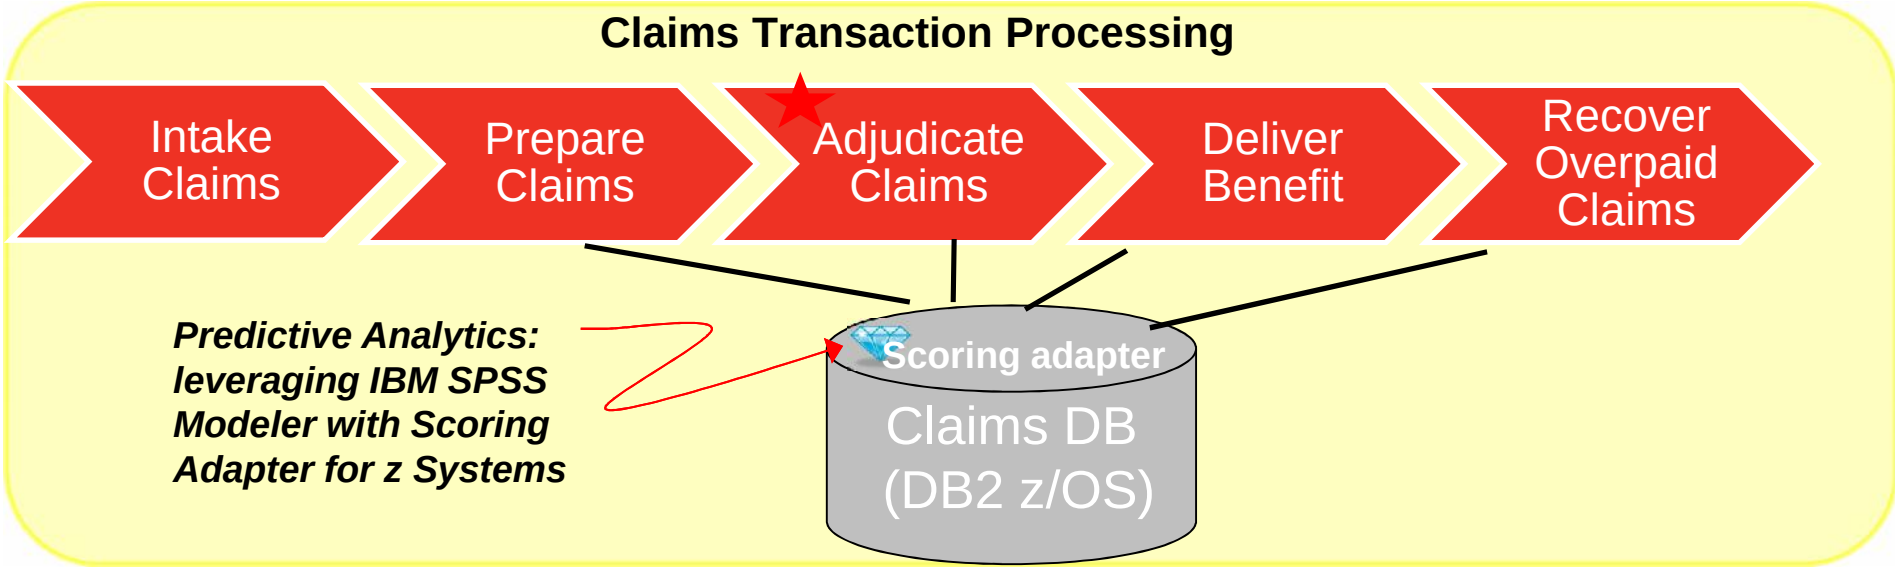

- Use predictive analytics to determine likelihood that a claim will be overpaid, appealed, etc. before the claim is paid, reduce pay-and-chase
- Use unique z Systems integration of real-time predictive analytics with DB2 z/OS to ensure the analytics will not impact claims throughput per day
- Demonstrated high levels of scale, low IT consumption and extremely fast real-time response times in customer engagements and internal demos
- Healthcare Insurance Demo using real-time scoring: http://youtu.be/_vII97Ylq0Y

Example: How pre-payment analytics can reduce costs

Pre-payment predictive models can:

- Target currently un-recovered over-payments
- Find potential over-payment prior to payment, avoiding cost and risk in recovery process
- Avoid post-payment recovery situations where legitimate submitters have accidently miscoded, helping to preserve valued relationships with clients

Base Assumptions – can vary

Claims per day	500,000
Claims auto-adjudicated	90%
Claims auto-paid per day	450,000
Avg payout per claim	$50
Total payout per day	$22,500,000
Incorrectly paid rate	1%
Claims incorrectly paid	4,500
Potential to be collected/saved per day	$225,000

Post-payment collection			Pre-payment predictive analytics	
Cost of collection per claim	$25	⟷	$0	Cost of collection per claim
Success rate	50%		90%	Success rate
Amount RECOVERED (50% of $225,000)	$112,500	⟷	$202,500	Amount NOT PAID (90% success rate out of $225,000)
Cost to recover claims ($25 per claim recovered)	$56,250		$0	Total cost to collect – overpayment not made
Total amount RECOVERED	$56,250	⟷	$202,500	Total amount NOT PAID

Potential savings per day $146,250

Potential savings per year $53Million+

Easy to deploy, simple to use Cloud Management Solution

New

Heterogeneous and integrated management support

- z Systems managing Power and x86 servers
- Central management across multiple hypervisors & domains
- All IBM server architectures & major hypervisors supported

Accelerated time to market with pattern support

- Chef-based patterns based on OpenStack Heat pattern engine is now supported on z Systems
- Workload deployment based on patterns speeds delivery of new services

Hybrid Cloud support

- Hybrid Clouds on and off premise options via SoftLayer support

openstack™

CHEF™

Quickly build out complex cloud workl instances on z Systems

New

- **Reduce deployment error/fix**
- **Reduce need for deep product skills**
- **Improve quality of delivery**
- **Reduces operating and capital expenses**

More

patterns to be delivered in 2015

2 patterns1

for key z System portfolio

- **WAS ND**
- **WAS Liberty**
- **ODM Decision Server**
- **ODM Decision Center**
- **Integration Bus**
- **DB2**
- **Business Process Server**
- **Business Process Center**
- **Business Monitor**
- **WebSphere Portal**
- **WebSphere MQ**
- **MobileFirst Platform**

Up to **80%** reduction in multi-product deployment

A private cloud on z13 yields the lowest TCO compared to a public cloud and a private cloud on x86

64 Light Workloads (128 VMs)

32 Medium Workloads (64 VMs)

27 Heavy I/O Workloads (27 VMs)

123 workloads (219 VMs)

Public Cloud

219 instances
$21.1M (3yr TCO)

Private Cloud on x86

264 x86 cores
$10.6 (3yr TCO)

Private Cloud on z13

32 IFLs
$7.2M (3yr TCO)

32% Less than x86 cloud*

66% Less than public cloud*

*estimated

Performance comparison based on IBM Internal tests comparing IBM z13 cloud with one comparably configured private x86 cloud and one comparably configured public cloud running an aggregation of light, medium and heavy workloads designed to replicate typical IBM customer workload usage in the marketplace. System configurations are based on equivalence ratios derived from IBM internal studies and are as follows: Public Cloud configuration: total of 219 instances (128 for light workloads, 64 for medium workloads and 27 for heavy workloads); x86 Cloud configuration: total of eleven x86 systems each with 24 Intel E7-8857 v2 3.0GHz cores, 512GB memory, and 7x400GB SSDs; z13 Cloud configuration: total of 32 IFLs, 3806GB memory, and Storwize v7000 with 47x400GB SSDs. Price comparison estimates based on a 3YR Total Cost of Ownership (TCO) using publicly available U.S. prices (including a 20% discount for middleware) current as of January 1, 2015. Public Cloud TCO estimate includes costs (US East Region) of infrastructure (instances, data out, storage, support, free tier/reserved tier discounts), middleware and labor. z13 and x86 TCO estimates include costs of infrastructure (system, memory, storage, virtualization, OS, cloud management), middleware, power, floor space and labor. Results may vary based on actual workloads, system configurations, customer applications, queries and other variables in a production environment and may produce different results. Users of this document should verify the applicable data for their specific environment.

A private cloud on z13 yields lowest TCO for a variety of workloads

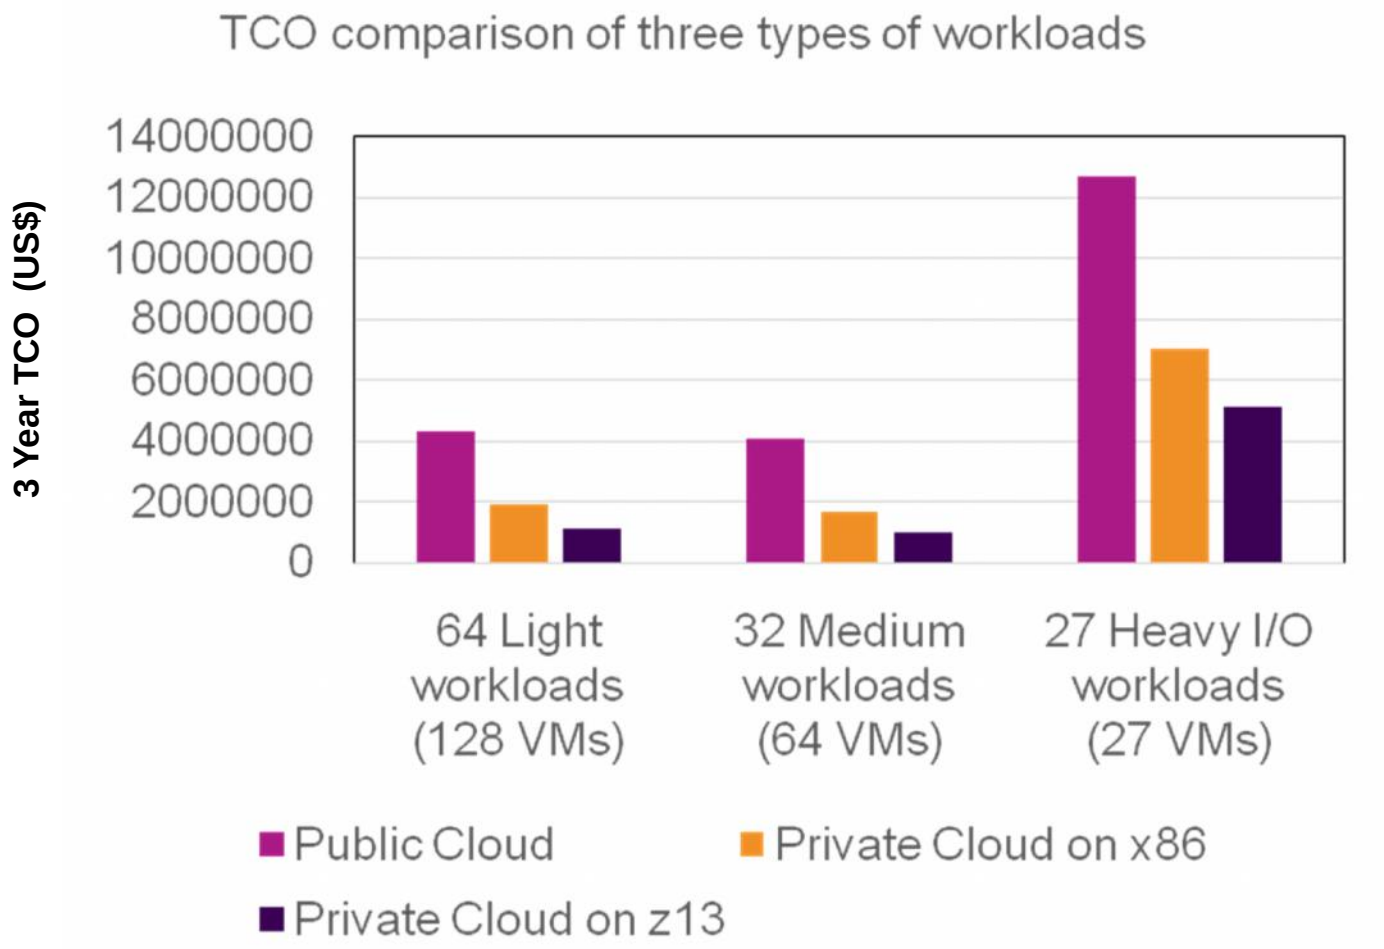

x86 and public cloud yield lower 3yr TCO

24 Web Workloads

24 VMs total

Compute-intensive web workloads
1 VMs per workload;
Each VM requiring 4GB memory; 20GB storage

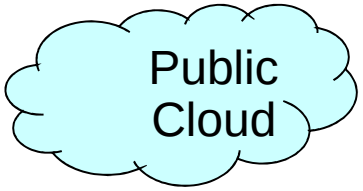

24 instances
$0.91M (3yr TCO)

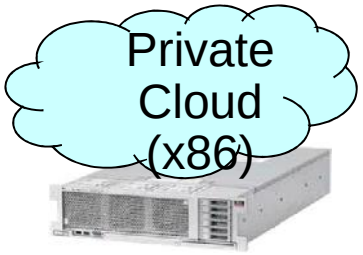

24 x86 cores
$0.61M (3yr TCO)

6 IFLs
$1.4M (3yr TCO)

System configurations are based on equivalence ratios derived from IBM internal studies.
Average utilization of 24-core x86 system is assumed to be 60%; avg utilization of z13 with 6 IFLs is assumed to be 75%; transaction response time is the same on all platforms

Create a competitive advantage in the digital era

Operational and analytics applications reside with the data in **a single fully virtualized system**

Creating **Intelligent interaction** in transactions integrated across mobile & core processes

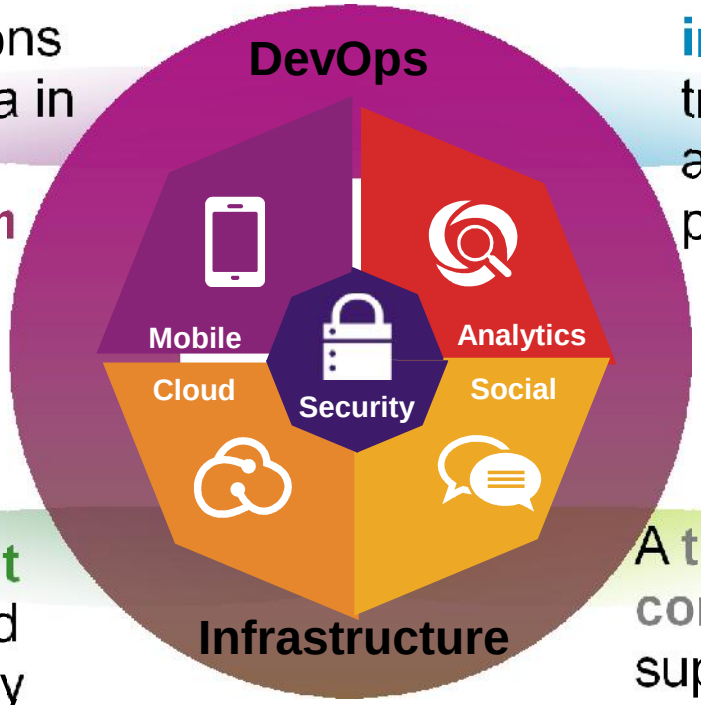

A **flexible and robust** infrastructure for rapid delivery of high quality services

A **trusted computing** platform to support secure growth in transactions and data

…make the extraordinary possible!!!

Thank You

The System z mobile development lifecycle

Design & Develop

Worklight Foundation
Rational Developer on System z

Instrument

Arxan Application Protection for IBM Solutions
Fiberlink MaaS360
Trusteer

Obtain Insight

Worklight Foundation
Tealeaf.
DB2 Analytics Accelerator

Integrate

Worklight Foundation
z/OS Connect
Cast Iron
DataPower
API Management
MQ/MQTT, MessageSight
IBM Integration Bus

Manage

Worklight Foundation
Fiberlink MaaS360
API Management

Test

Rational Test Workbench
Worklight Mobile Quality Assurance

Has a z Systems part number

Deploy

Worklight Foundation
UrbanCode Deploy

Scan & Certify

AppScan including Worklight Application Scanning


You know? you can do this with your mobile device now.
P.C.VEY

Securing mobile apps for System z

1 **Unified management and security control for mobile devices**
Fiberlink MaaS360
Trusteer

2 **Enforce web and mobile security policies**
IBM DataPower
IBM CastIron

3 **Protect against fraudulent and unauthorized access**
IBM Security Access Manager for Mobile
Arxan Application Protection for IBM Solutions

4 **Security foundation for System z mobile workloads**
IBM Security zSecure Suite, RACF, LDAP,
IBM Security QRadar SIEM,
IBM InfoSphere Guardium