17079: The Economics of Mobiles and Mainframe

David Rhoderick, IBM

#SHAREorg

SHARE is an independent volunteer-run information technology association that provides **education**, professional **networking** and industry **influence**.

Agenda

- Mobile revolution is trending more than *“eBusiness 1.0”*
- MobileFirst deployment choices
- Mobile Workload Pricing
- Architectural choices for data serving

Mobile business is a much bigger phenomenon than "ebusiness 1.0"

Time Magazine, January 2014

Projections of mobile growth and PC decline based on Gartner data

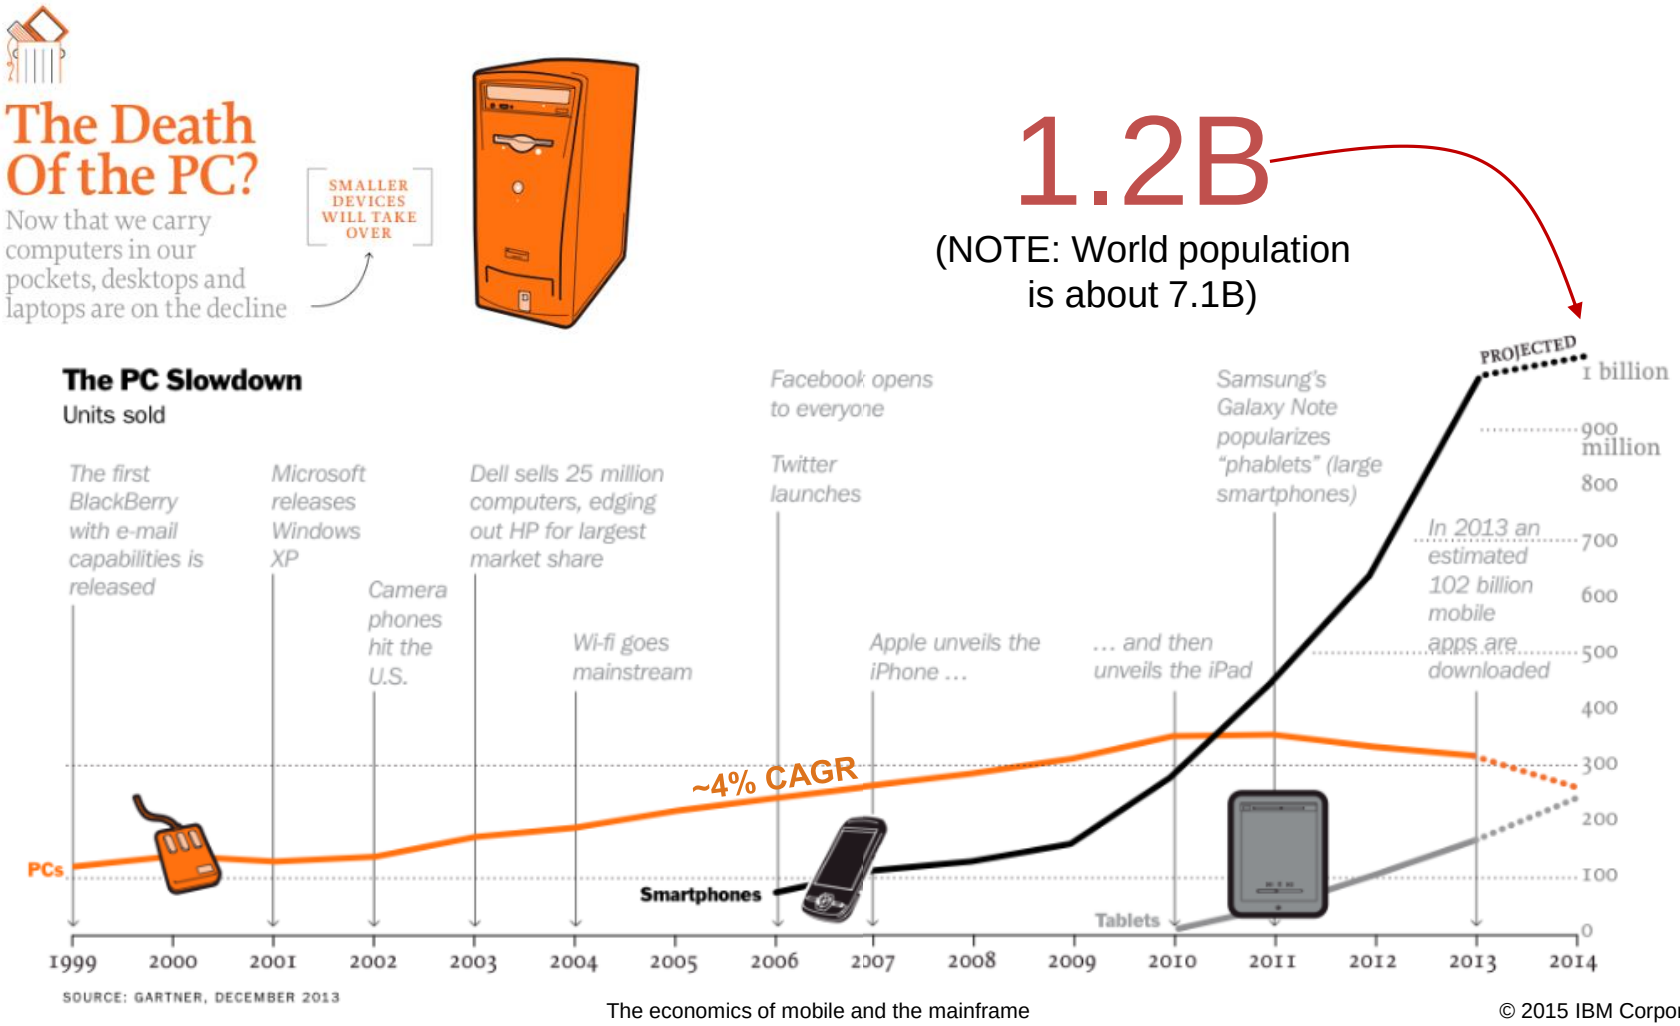

Installed smart devices probably now exceed PCs (trailing 5y PC sales and trailing 24m phone/tablet sales)

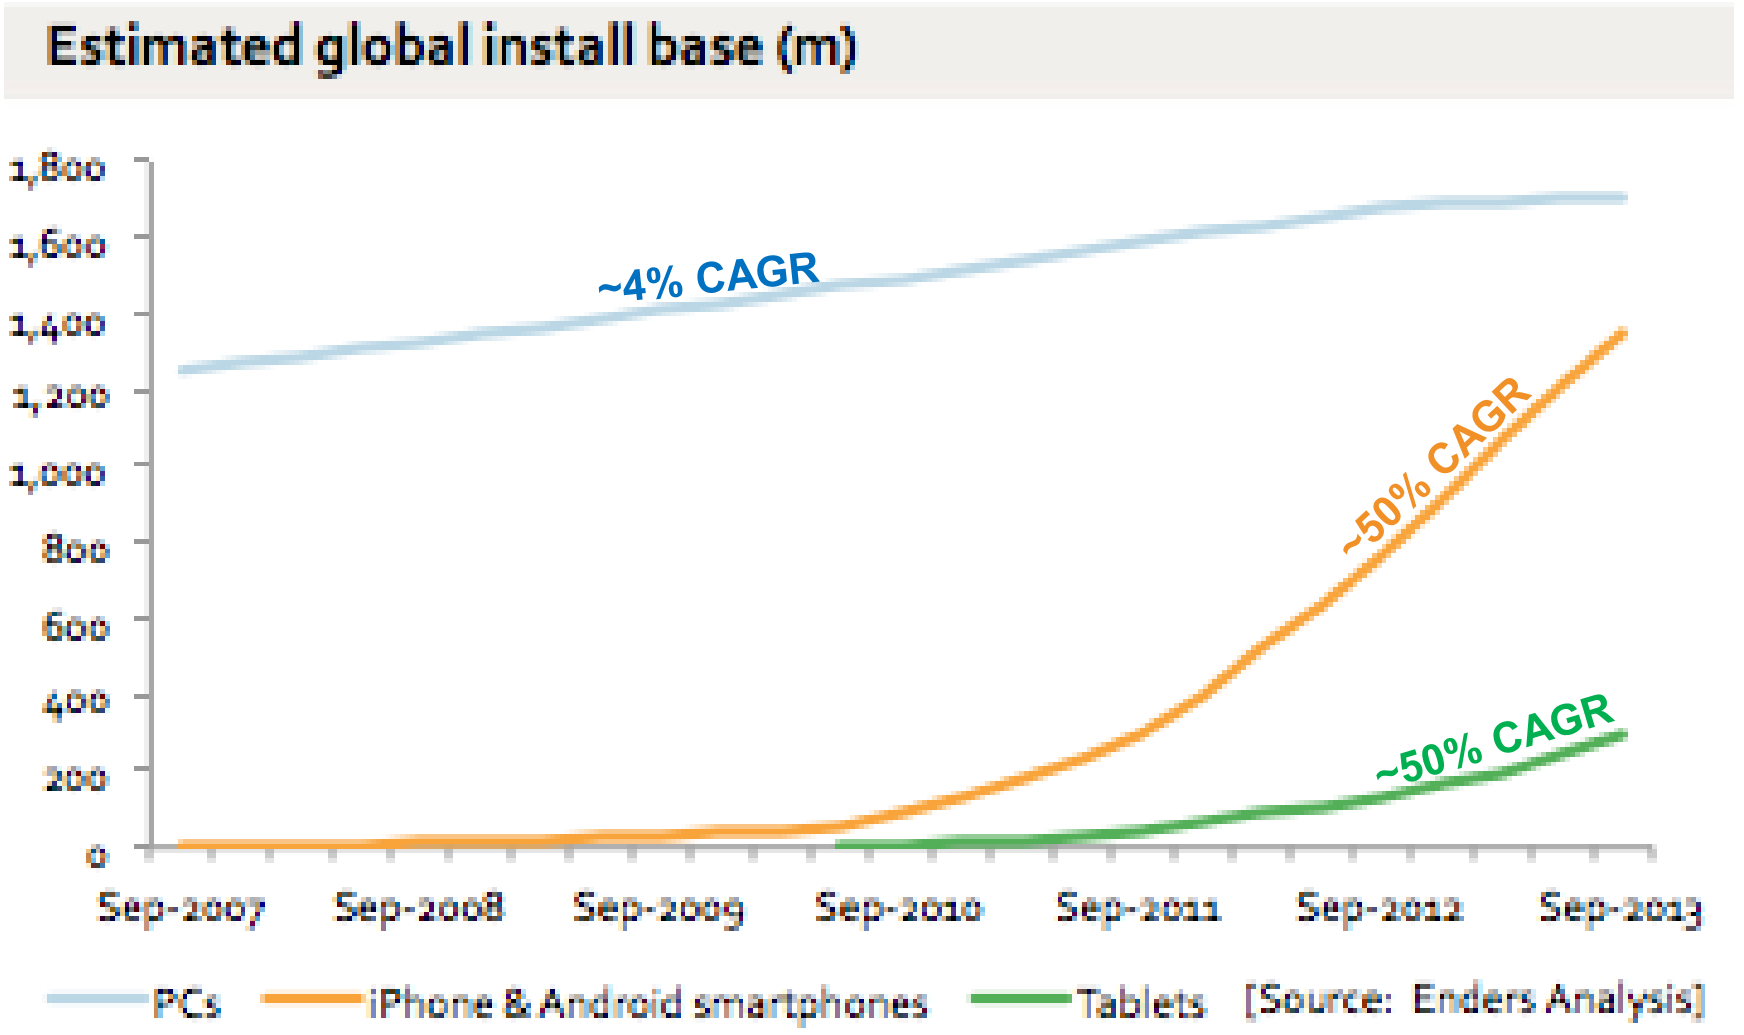

Other signs that the technology is becoming widely adopted

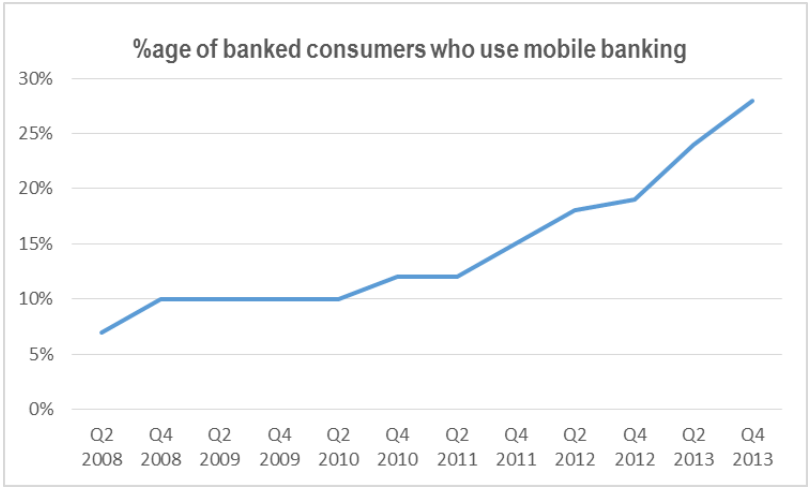

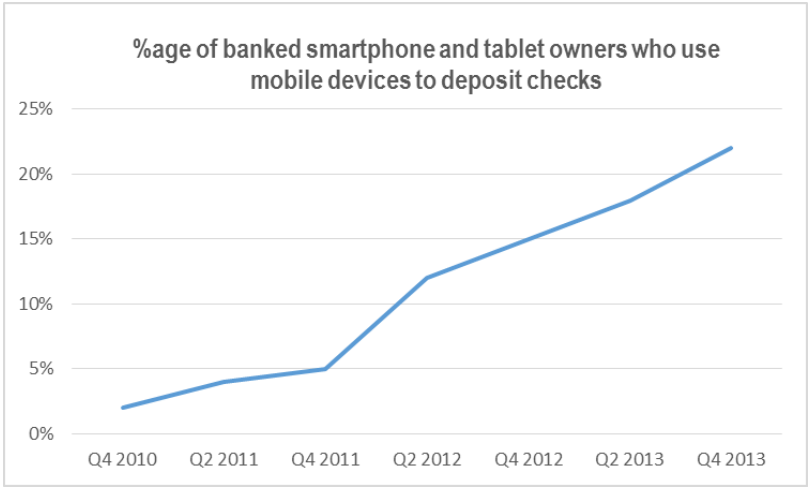

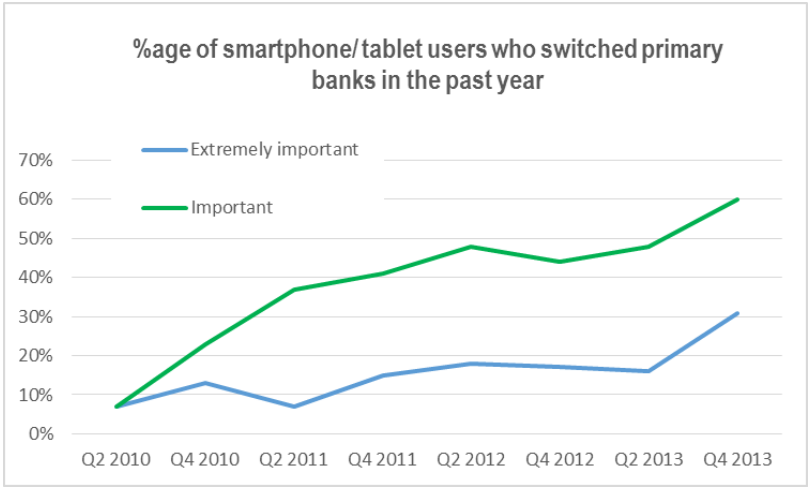

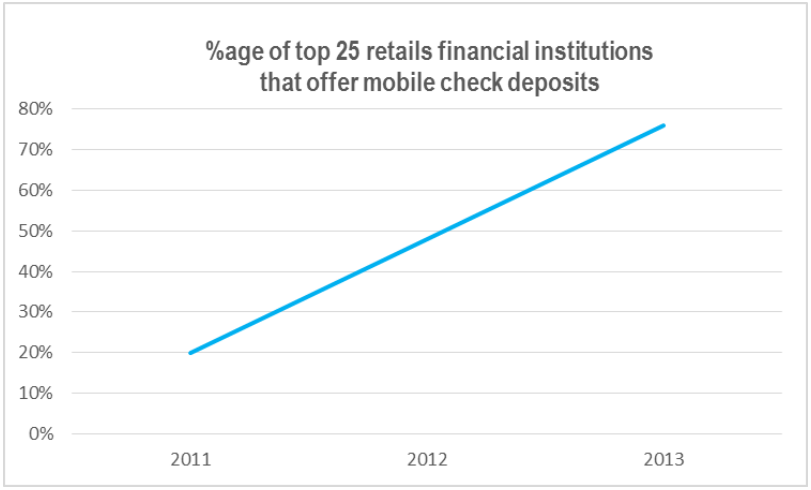

Sources: Javelin Strategy and Research, Alix Partners
See http://online.wsj.com/news/interactive/mobile0326A?ref=SB10001424052702303847804579481811781070456

Mobile Traffic as % of Global Internet Traffic = Growing >1.5x per Year & Likely to Maintain Trajectory or Accelerate

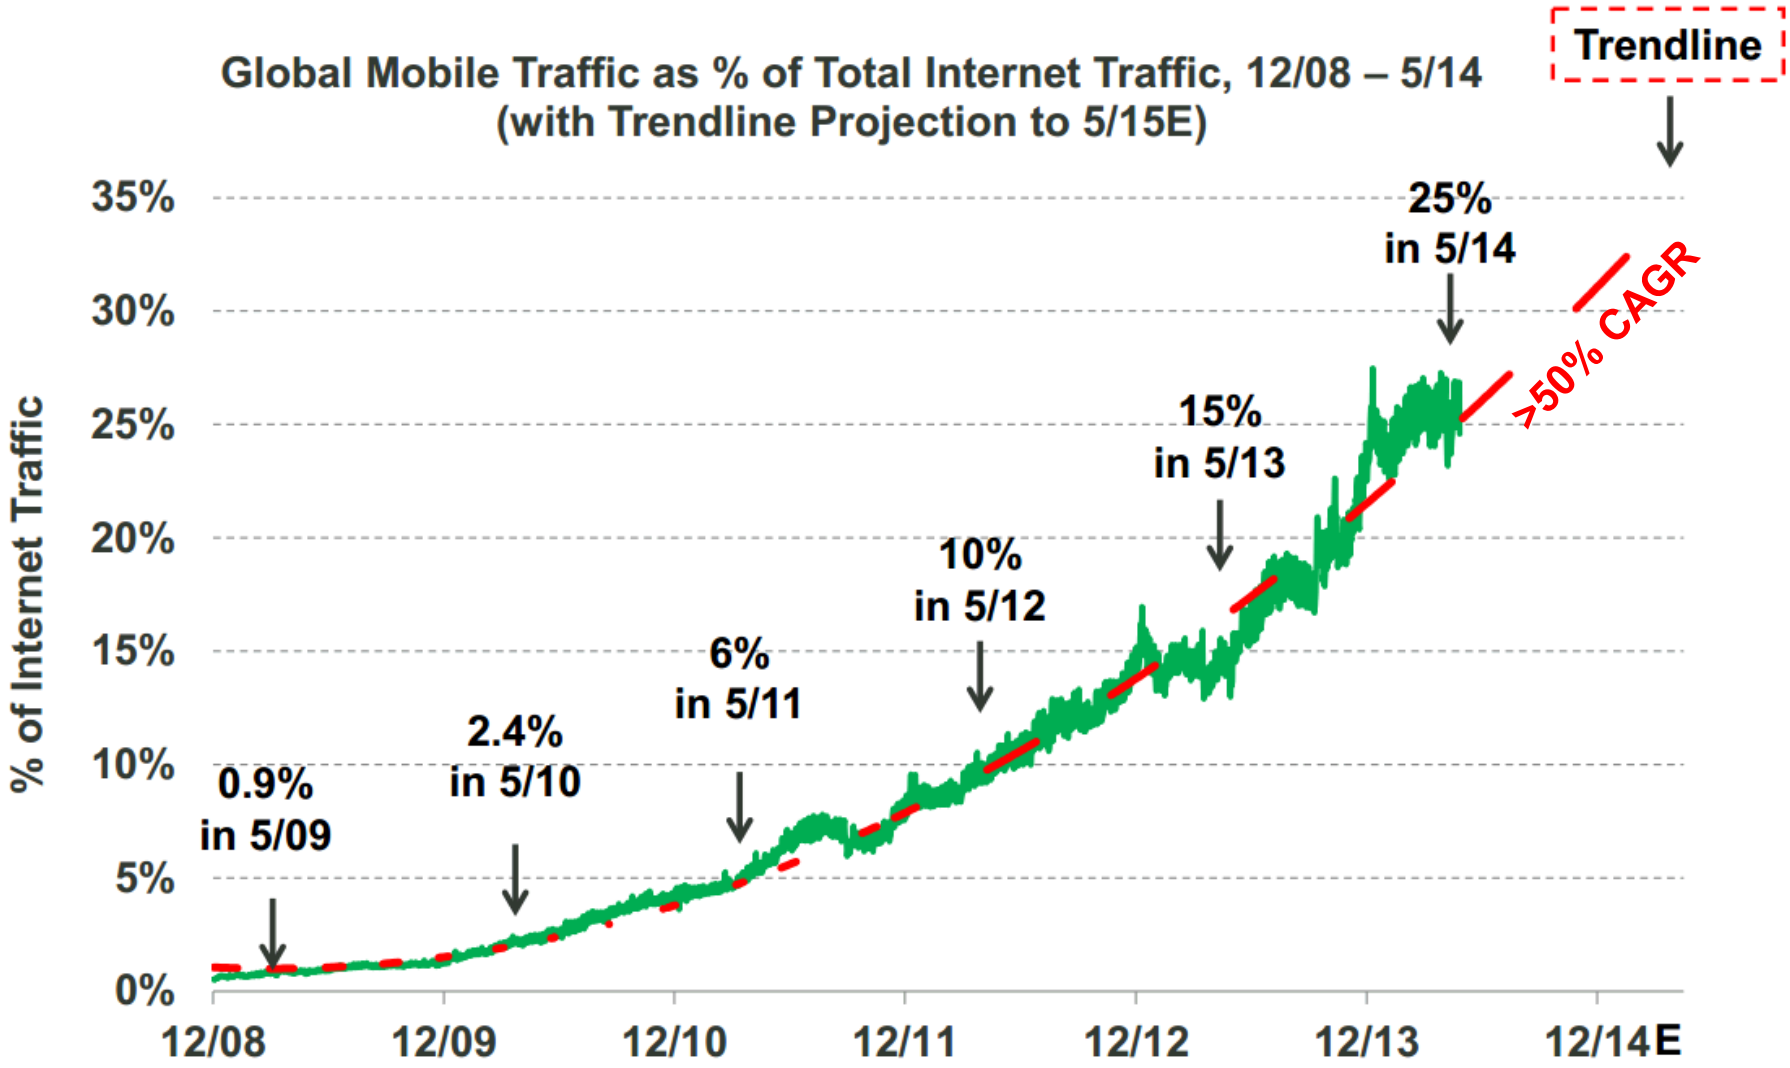

@KPCB

Source: StatCounter Global Stats, 5/14. Note that PC-based Internet data bolstered by streaming.

Source: Mary Meeker, 5/28/2014

FORECAST: US E-Commerce, M-Commerce And Mobile In-Store Payments

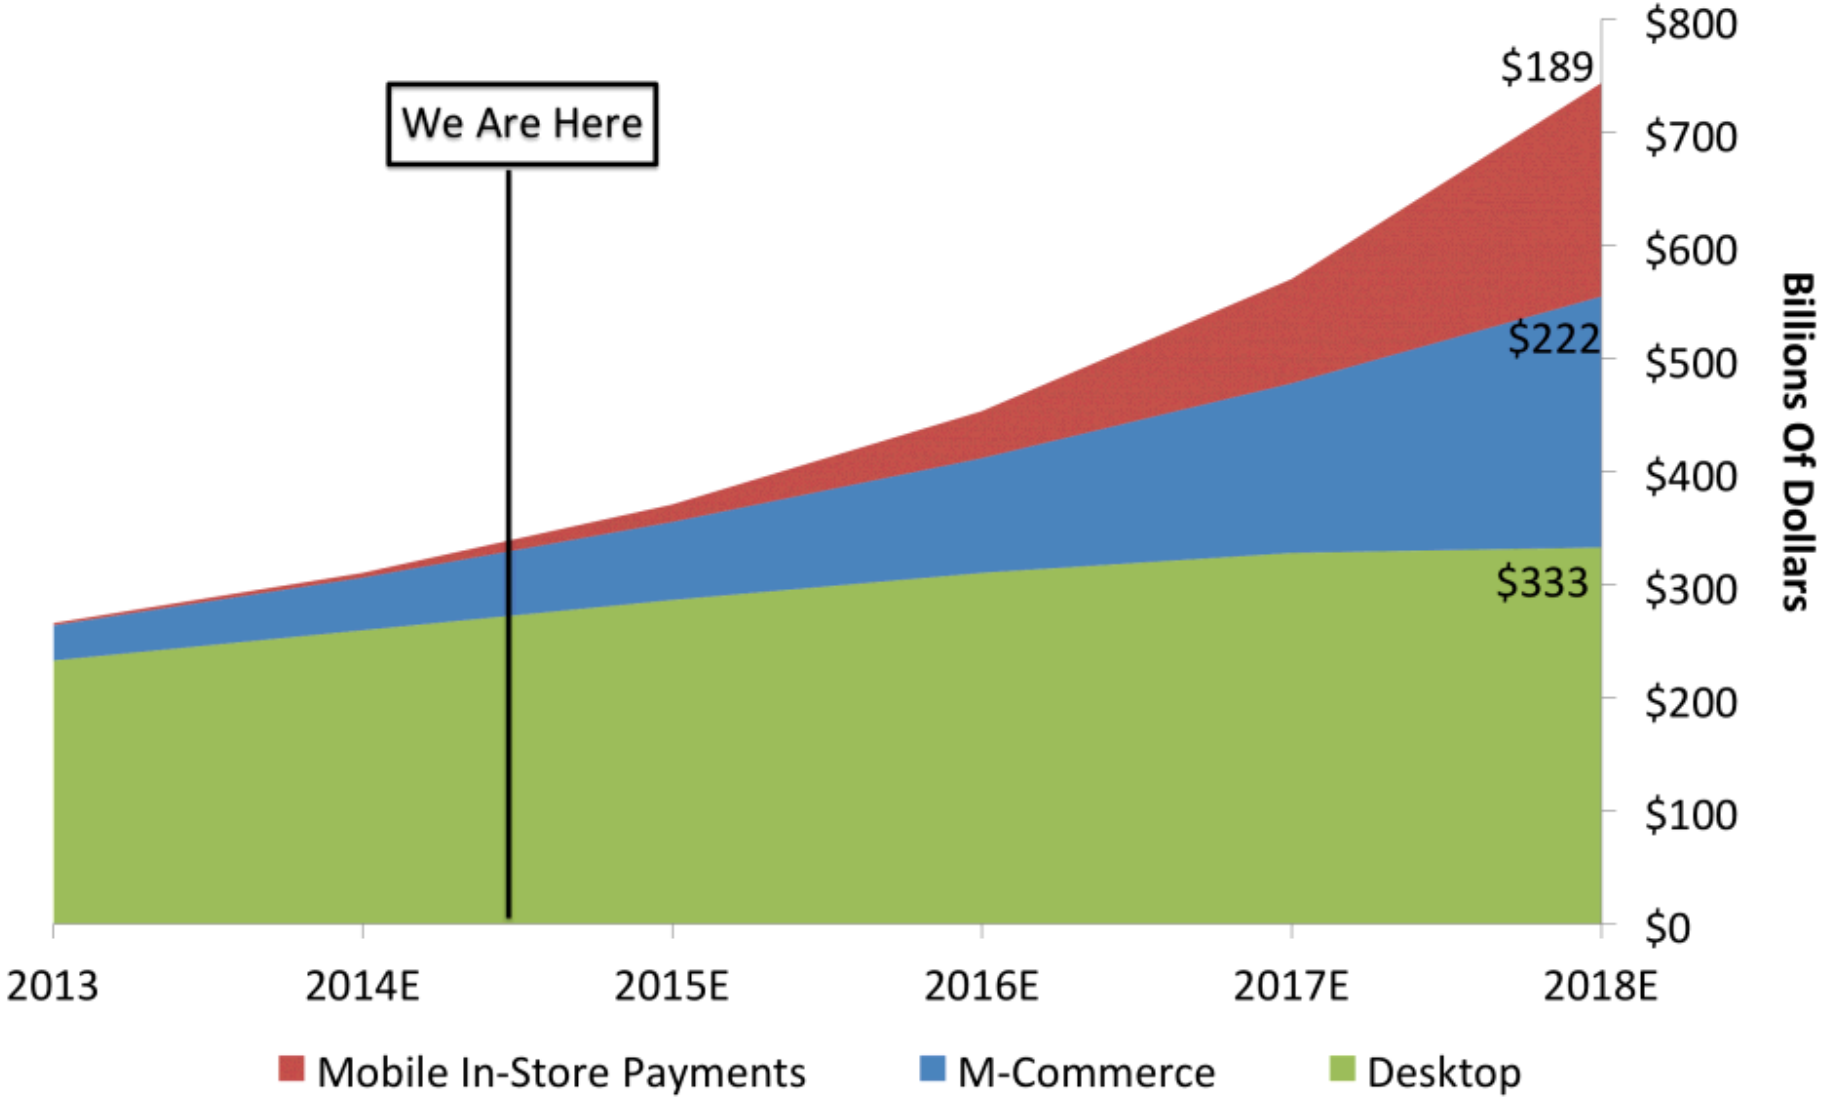

Source: BI Intelligence

BI INTELLIGENCE

Web-enablement of back-end functions is effective way to grow business – “ebusiness” includes modernizing existing assets

- In September 2012, the independent research firm Vanson Bourne studied IT trends and interviewed 520 CIOs from large enterprises across a range of industries in the U.S., Europe & Asia

- ***62% have enabled back-end functions to support e-commerce and mobile applications***
  - Mature technology and tooling helps do this quickly and easily

- ***68% also believe mobile computing will drive MIPS growth***

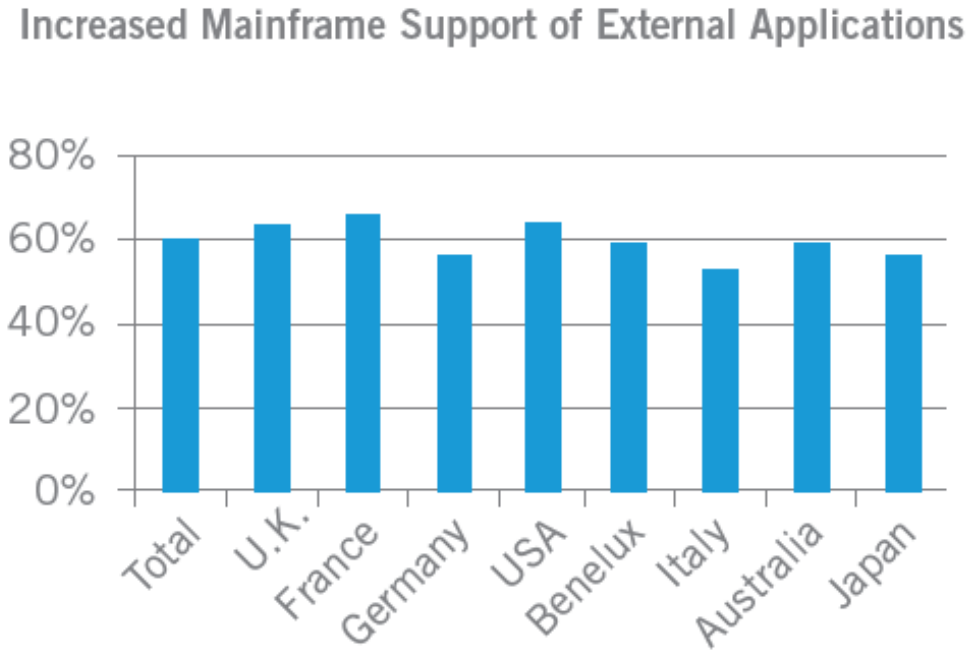

System z installed capacity grown at 17% CAGR – partly due to "eBusiness 1.0" – direct mobile access will accelerate this

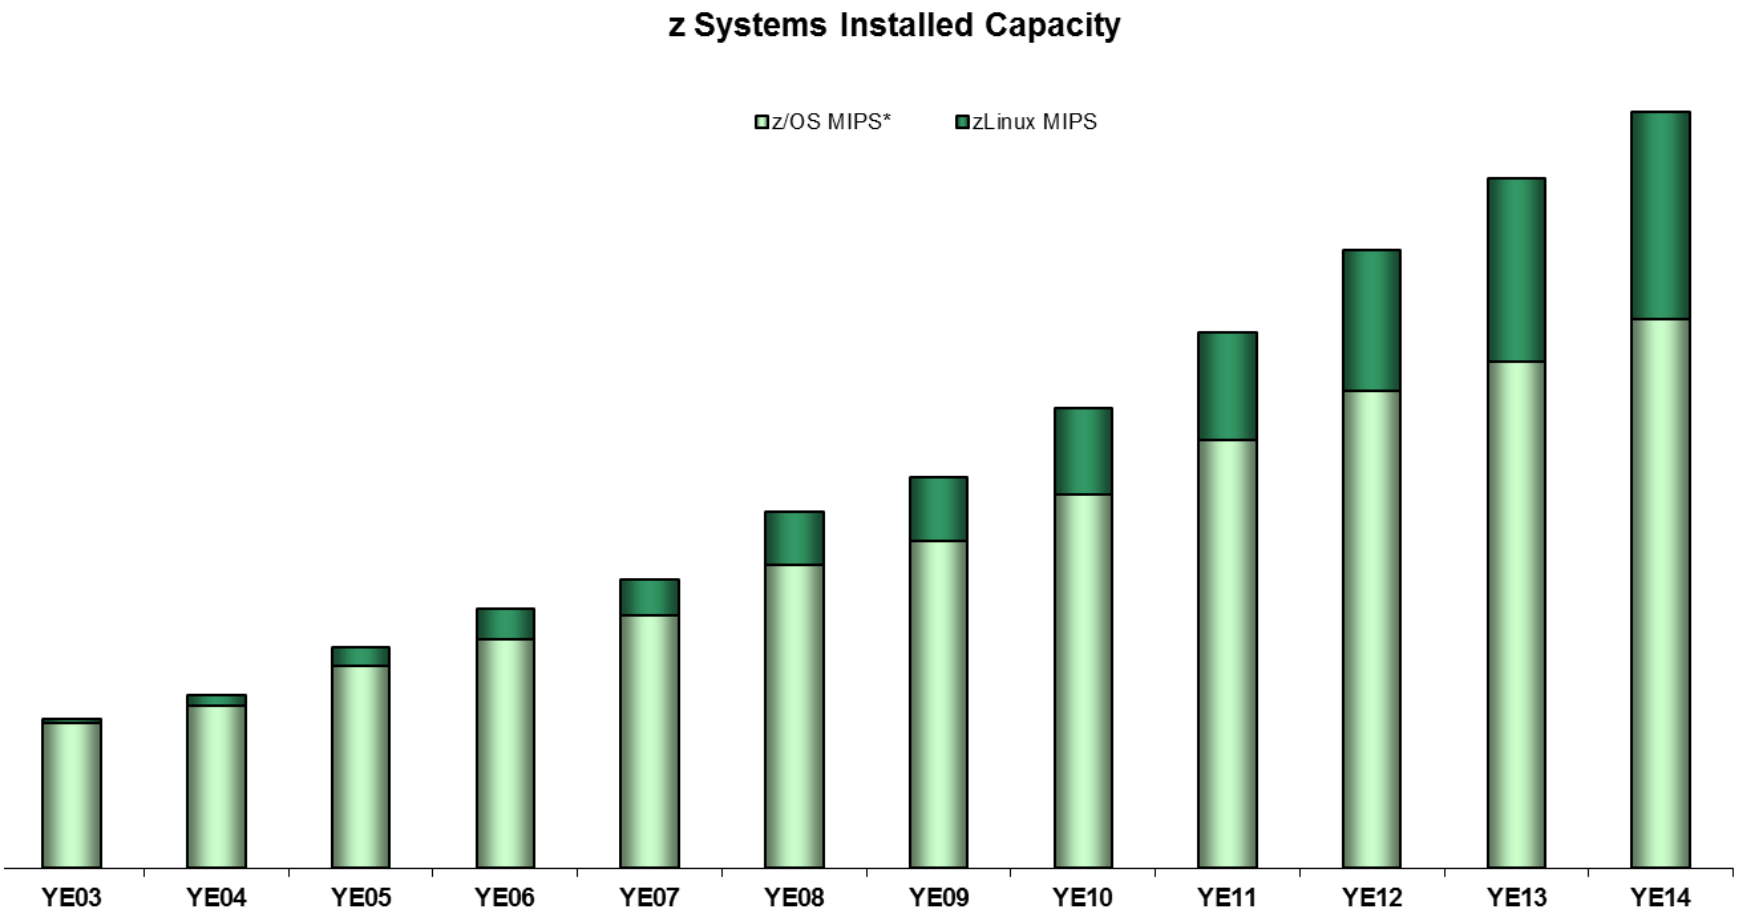

** these MIPS could also be used for z/VM, z/VSE or z/TPF Operating Systems*

There are proactive – and reactive – reasons for opening a mobile channel

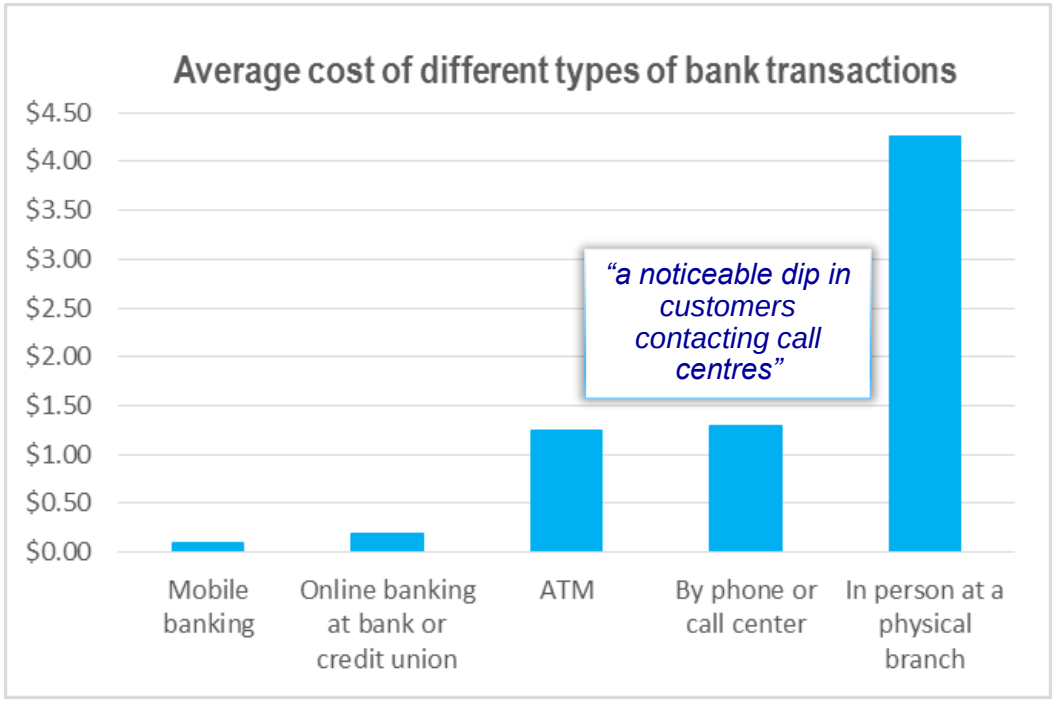

"… if we don't offer the functionality desired by our customers, there's **significant risk of falling behind** in the market where competitors are already offering these services and we will **fail** to attract sufficient custom to achieve the target **to grow market share**

"… Smartphone channels offer a true alternative in many instances to the internet self-service channels…"

Sources: Javelin Strategy and Research, Alix Partners
See http://online.wsj.com/news/interactive/mobile0326A?ref=SB10001424052702303847804579481811781070456

Mobile is changing the way customers interact – far more possibilities

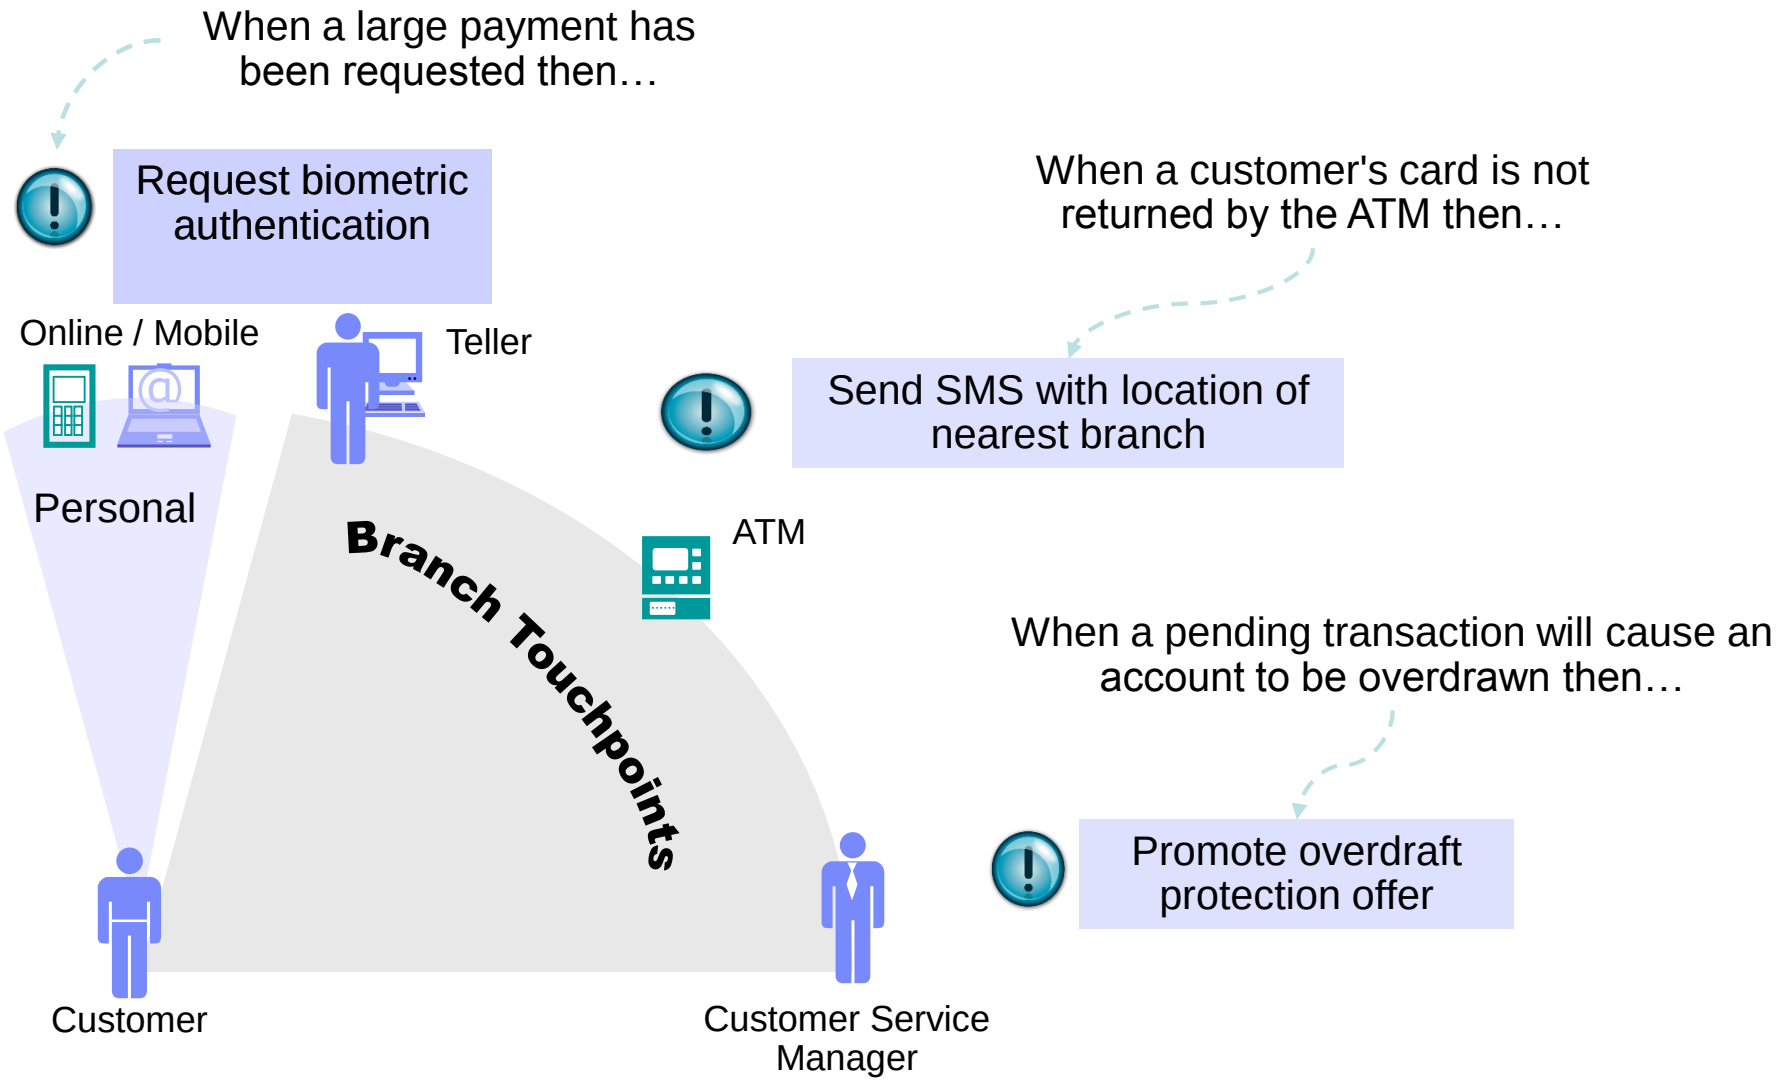

IBM z/OS Connect mobile enables backend on System z

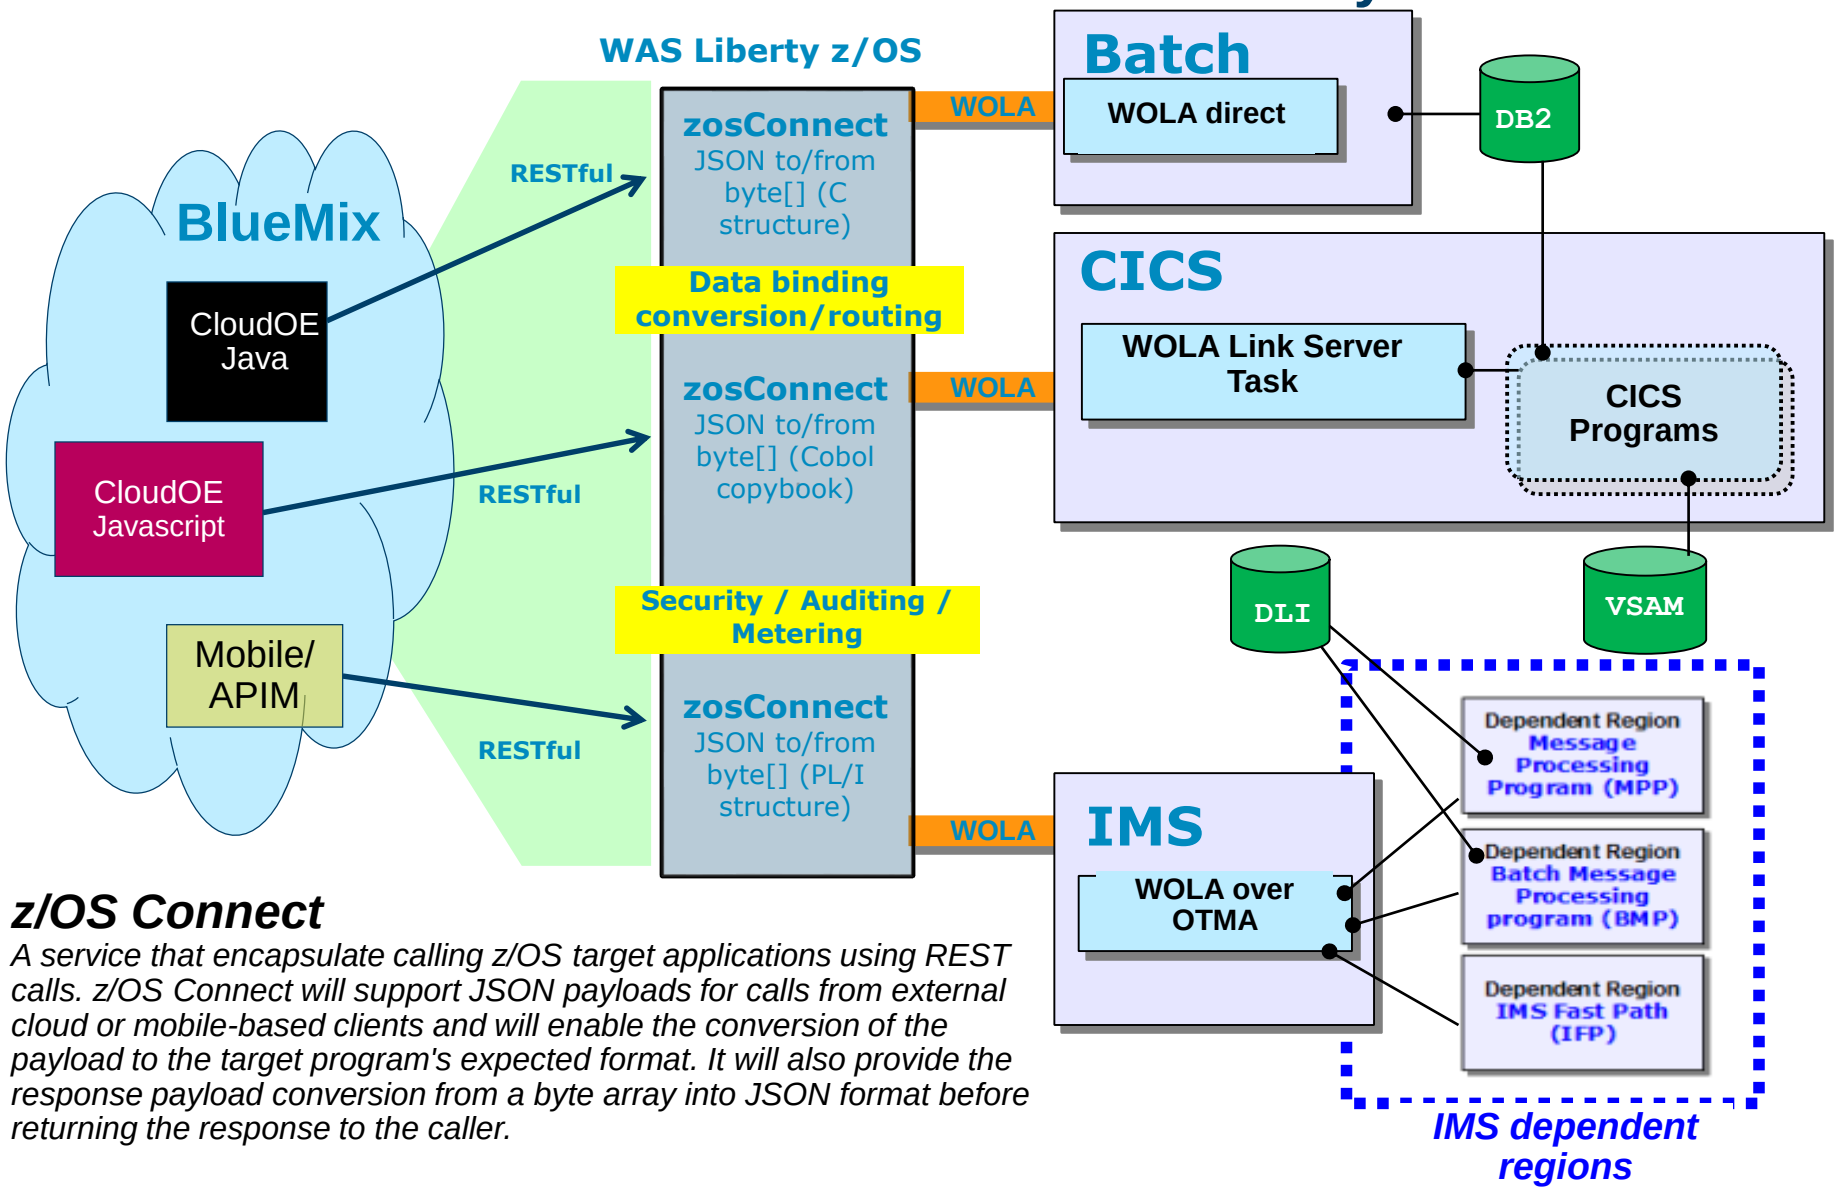

z/OS Connect

A service that encapsulate calling z/OS target applications using REST calls. z/OS Connect will support JSON payloads for calls from external cloud or mobile-based clients and will enable the conversion of the payload to the target program's expected format. It will also provide the response payload conversion from a byte array into JSON format before returning the response to the caller.

Centralized server technology provides a platform to manage and drive all mobile applications

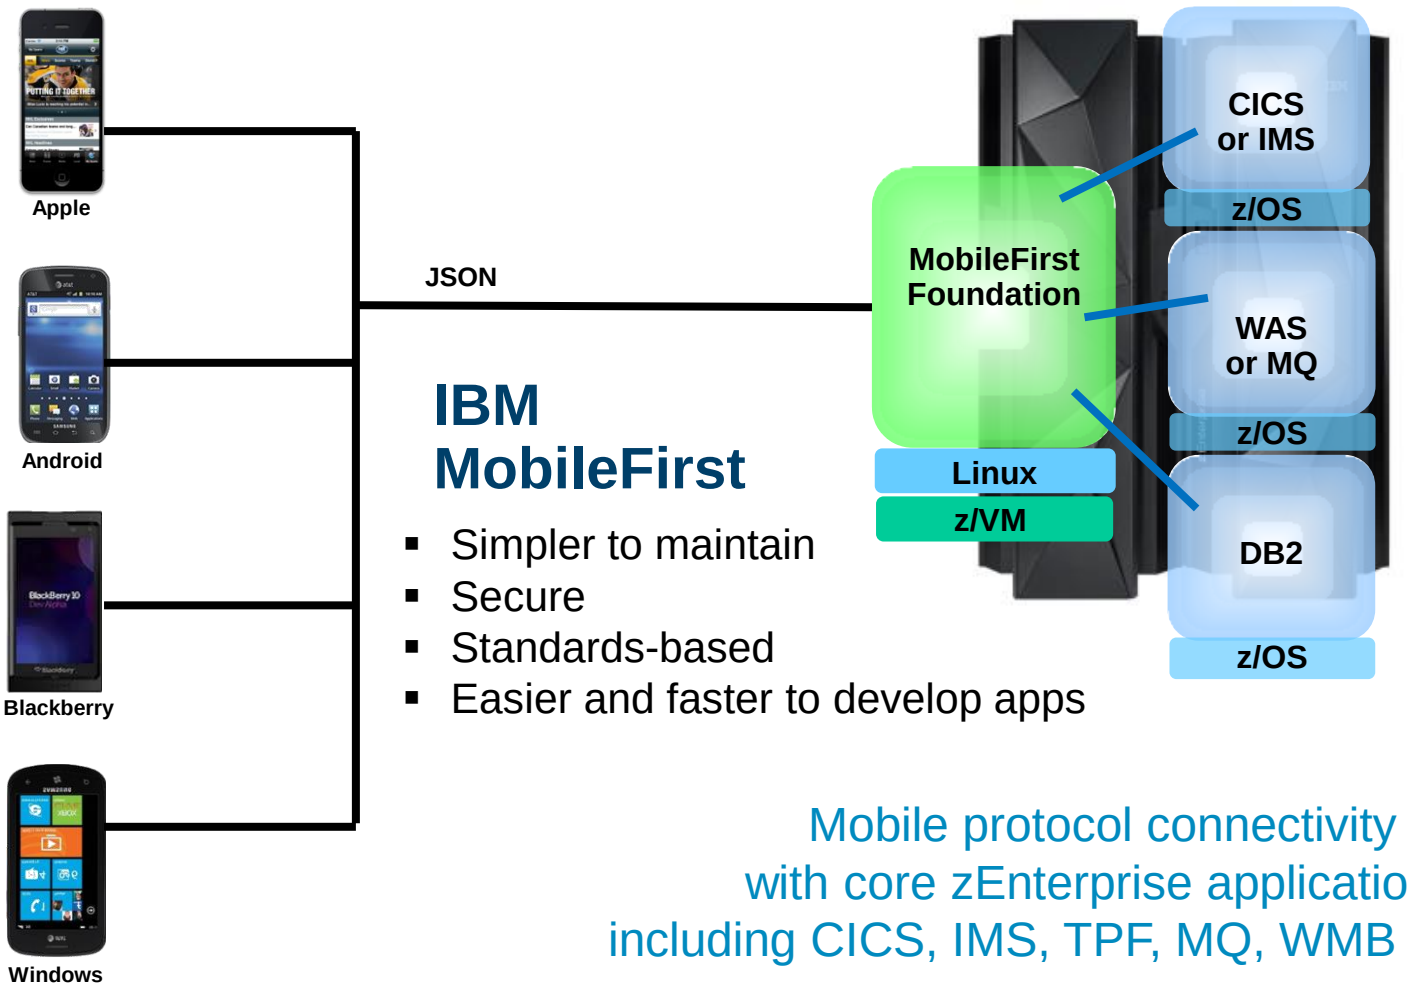

IBM MobileFirst

- Simpler to maintain
- Secure
- Standards-based
- Easier and faster to develop apps

Mobile protocol connectivity
with core zEnterprise applications
including CICS, IMS, TPF, MQ, WMB and DB2

DevOps capabilities for mobile applications with enterprise systems

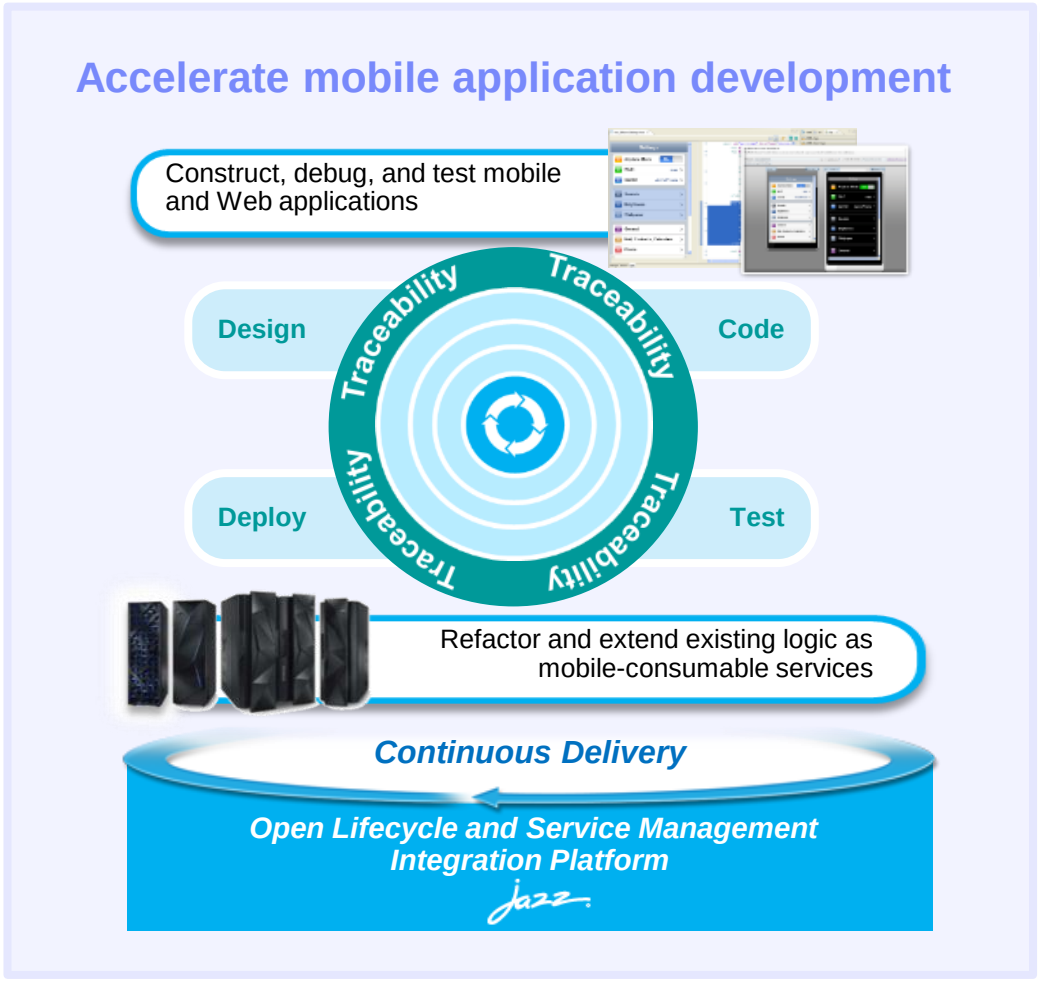

Rational Test Workbench
IBM MobileFirst Platform
IBM DevOps services for Bluemix
Rational Application Developer
IBM UrbanCode Deploy
Rational Developer for the Enterprise
CICS Interdependency Analyzer
CICS Performance Analyzer

Enabling you to…

- Quickly design, code, build, test, and deploy mobile apps that run on a wide variety of mobile platforms
- Refactor and extend existing back-end services to provide an optimal mobile experience
- Stub out back-ends to simplify test and development
- Automate testing for native and hybrid mobile apps

University of Florida goes mobile

UF UNIVERSITY of FLORIDA

Enabling 50,000 students, 5,400 faculty members and staff access to online features anytime, anywhere

Data provided to students real time
Mobile formatted information of class schedules, textbooks, academic dates, grades, emergency information and campus map

IBM Solution
Accessing CICS with System z information via smartphones

Up to 1M transactions/day

HALKBANK

25 years of continuous operation
On CICS®, DB2®, and System z®

Mobile

"*Mobile banking* is a fast growing channel for Halkbank. As for all other channels, *we run the transactions themselves on our CICS and DB2 for z/OS systems*, which helps *to ensure the highest levels of performance, availability and security*"

Ayhan Yalkut, System z Manager, Halkbank.

Scalable

The Turkish state pays most of the government salaries through the bank, causing a once-monthly peak in daily transactions, increasing from 25m to 50m

Reliable

"In more than 25 years of continuous operations, we have never had any unplanned downtime"

FORRESTER®

"Core enterprise applications process transactions — lots of transactions — and so can't be scaled by adding load balancers and spawning new copies of services. They scale by adding capacity to the transaction processing engine"

IBM studied MobileFirst Platform running on Linux/x86 versus Linux/z

- CICS/DB2 Application on System of Records (GENAPP)
  - Montpellier Benchmarking System Center
- Tested 400 concurrent threads with 60% Login, 30% Add or Delete and 10% Update
- Compared Throughput, Response Time and TCA per TPS between Linux on System z versus x86
  - 61% higher throughput
  - 36% lower response time
  - 10% lower TCA

Architecture of MobileFirst Platform Performance Test

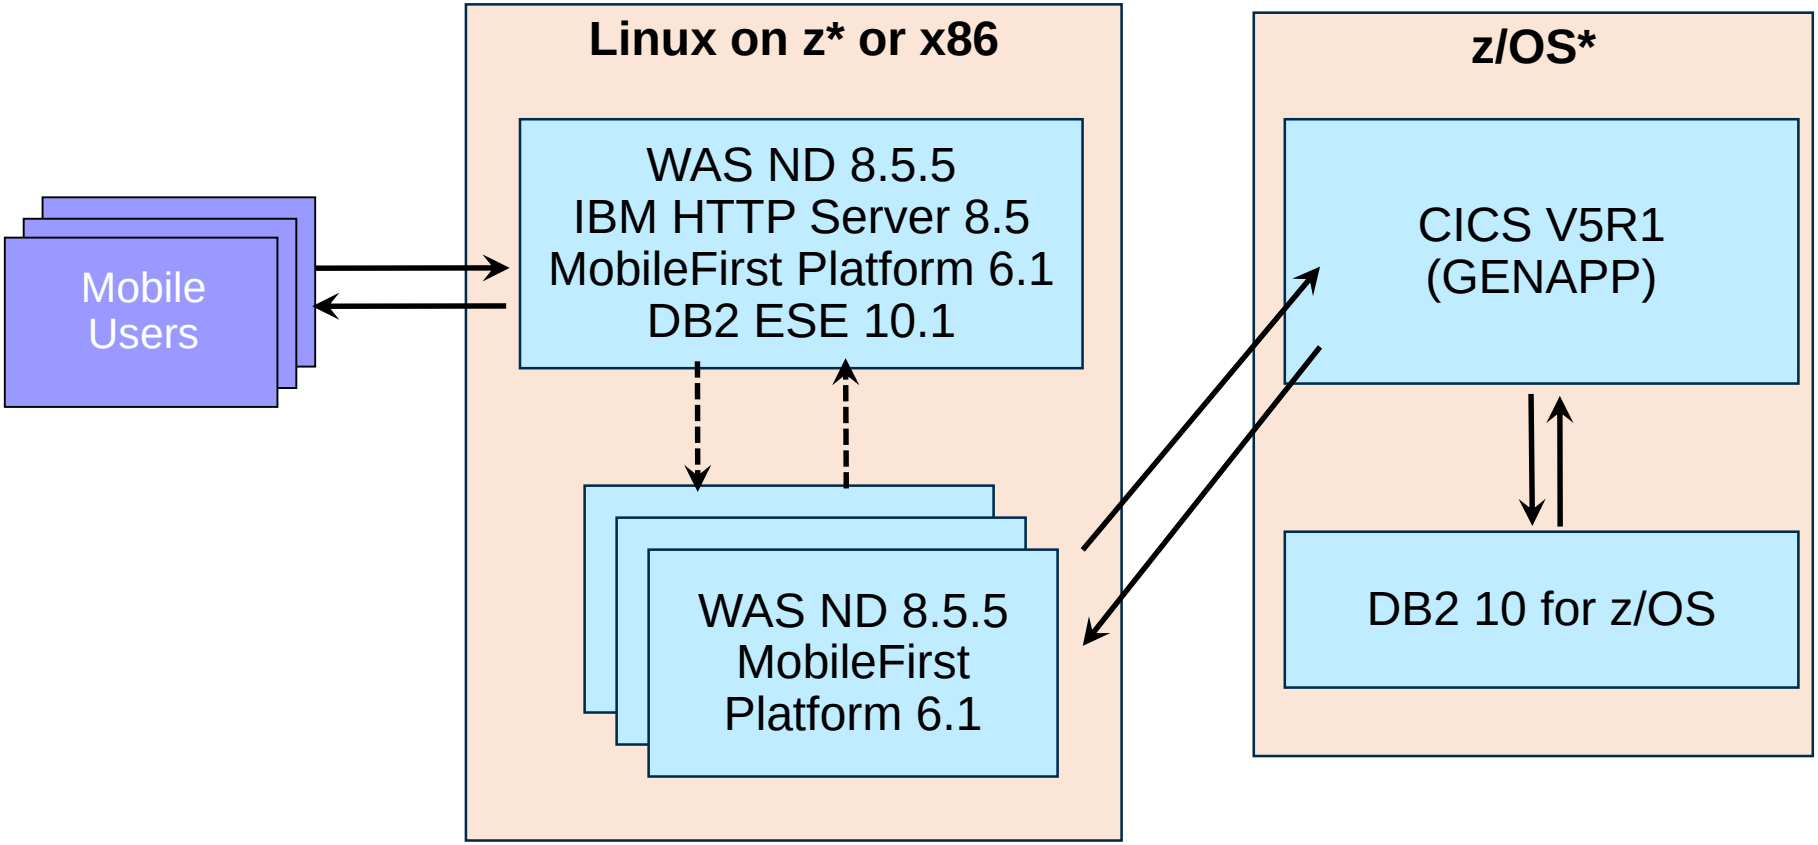

* Estimated performance, sizing and cost for z13 based on tests conducted on zEC12

This is based on an IBM internal study designed to replicate a typical IBM customer workload usage in the marketplace. Test involved 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 published US list 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.

Raw test results summary

- Test 4: Mix (60% Login, 30% Add or Delete, 10% Update)

		Worklight server on z/Linux			Worklight server on x86		
Number of CPU per Worklight server	Number of users	Response Time (ms)	Throughput (TPS)	Max number of physical CPU used (on 8CPs)	Response Time (ms)	Throughput (TPS)	Max number of physical CPU used (on 16Cores)
1	10	42.9	295.71	0.56 (avg 0.55)	50.2	242.89	0.75 (avg0.68)
	30	44.5	850.4	2.17 (avg 1.90)	54.9	652.6	1.7 (avg1.7)
	50	47.7	1304.1	2.32 (avg 2.21)	59.5	991.8	2.5 (avg2.14)
	100	60.3	2000	3.30 (avg 3.18)	75.4	1518.1	3.87 (avg3.87)
	200	104.4	2195.8	3.67 (avg 3.41)	132.2	1676.2	4.26 (avg4.26)
	400	189.9	2361.2	3.88 (avg 3.78)	256.8	1699.1	4.64 (avg4.64)
	600	293.1	2248.1	3.96 (avg 3.71)	357.8	1843.8	4.7 (avg4.7)
2	10	42.8	296.34	0.61 (avg 0.59)	50	244.12	0.75 (avg0.75)
	30	43	881	2.3 (avg 1.74)	54.4	659	1.83 (avg1.83)
	50	44.4	1418.9	2.48 (avg 2.42)	62.5	931.8	2.3 (avg2.3)
	100	48.6	2554.3	4.30 (avg 4.21)	68.7	1675	4.14 (avg4.14)
	200	70.8	3320.2	6.02 (avg 5.34)	95.7	2337.7	6.08 (avg6.08)
	400	131.4	3446.1	5.96 (avg 5.53)	205.4	2145.2	4.96 (avg4.96)
	600	215.1	3126.9	5.97 (avg 5.08)	291.6	2264.7	6.05 (avg6.05)

MobileFirst Platform scales better on Linux on System z than x86

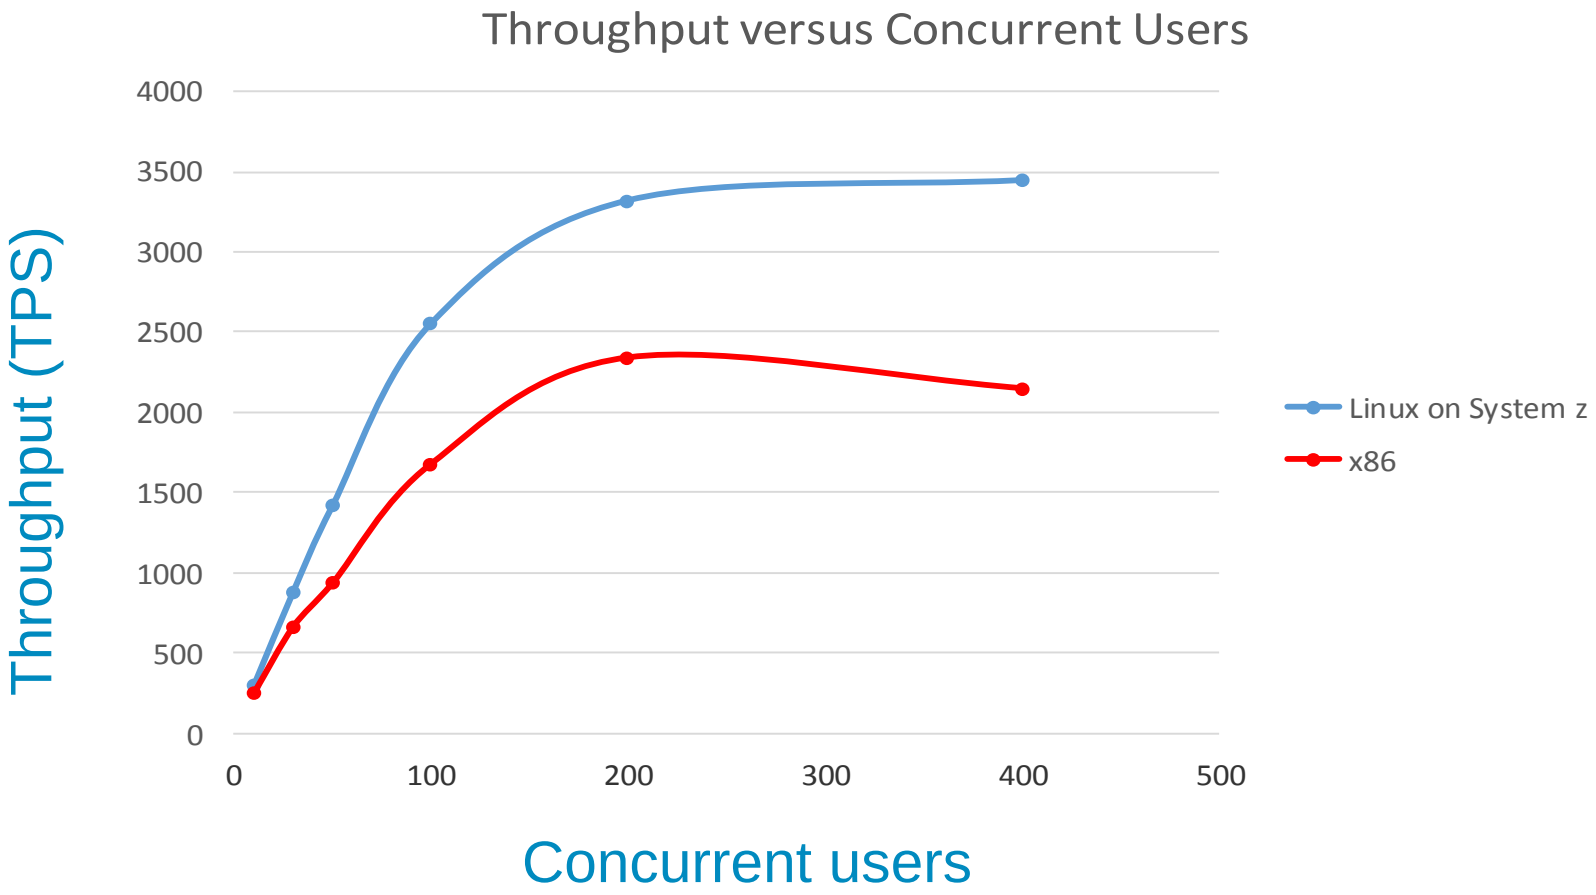

* Estimated performance, sizing and cost for z13 based on tests conducted on zEC12

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.

MobileFirst Platform provides lower Response Time on Linux on z than x86

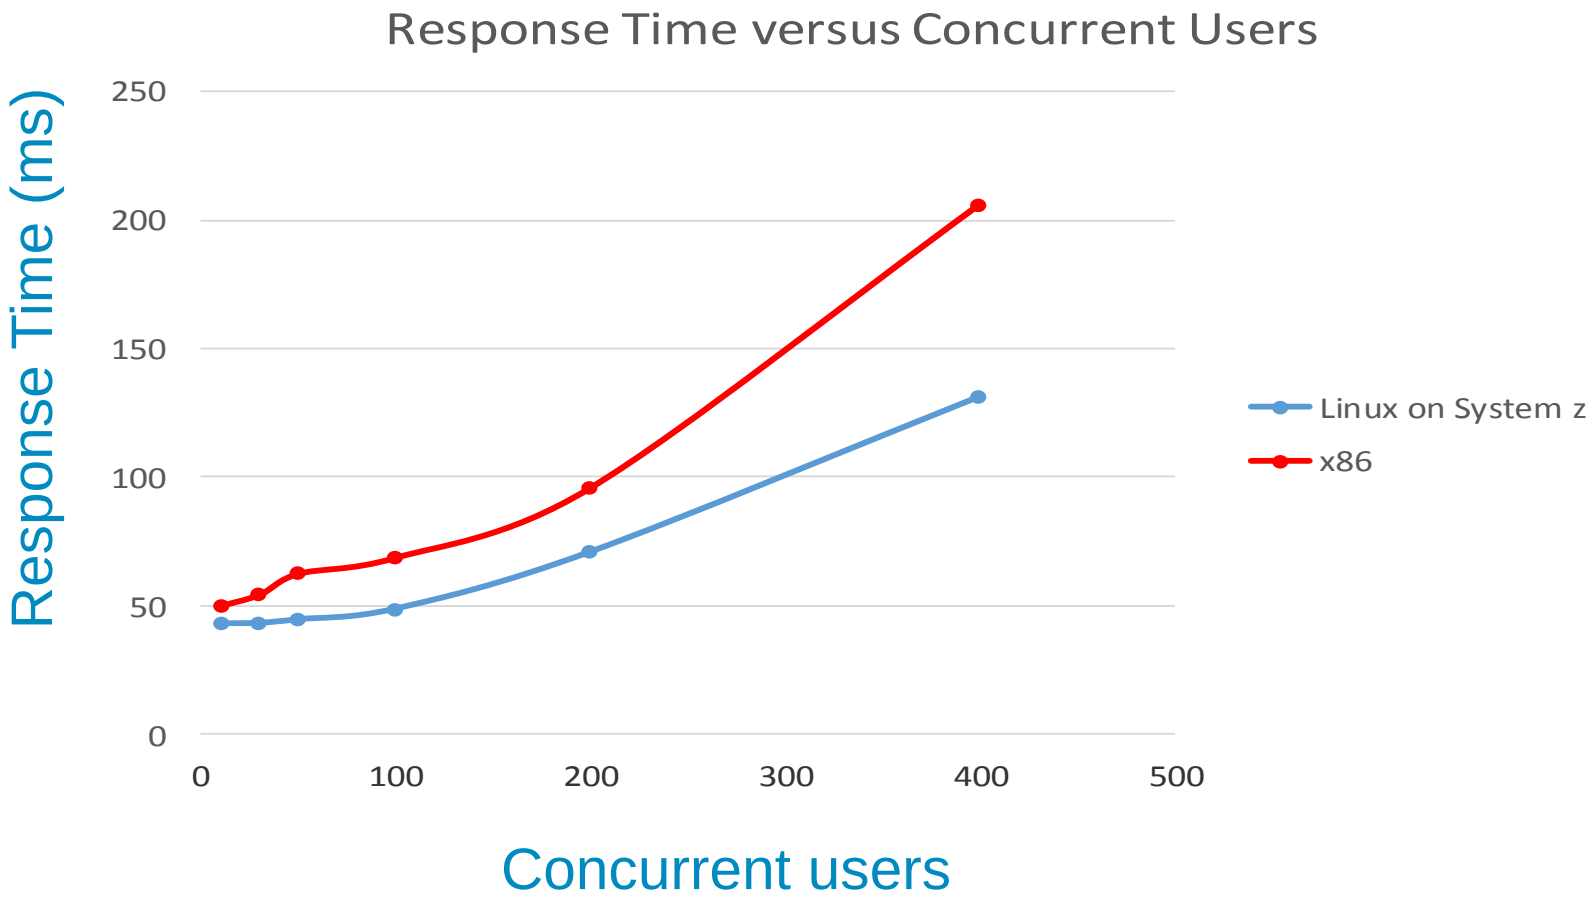

* Estimated performance, sizing and cost for z13 based on tests conducted on zEC12

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.

MobileFirst Platform is expected to provide lower front-end cost on Linux on z* than x86, for > 30 concurrent users

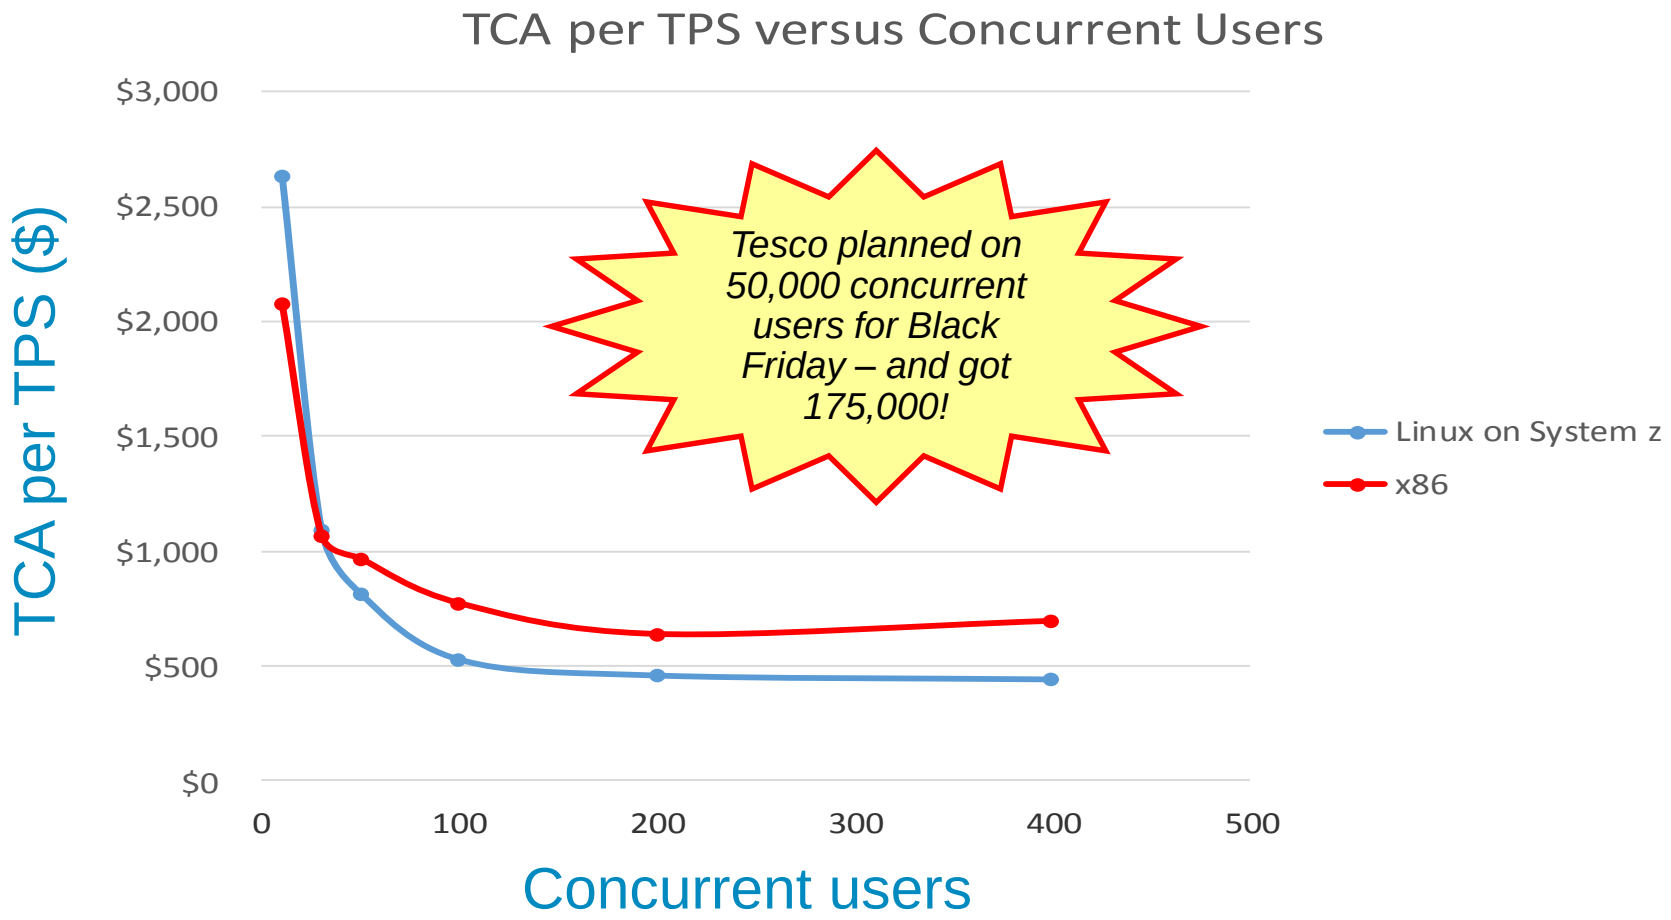

* Estimated performance, sizing and cost for z13 based on tests conducted on zEC12

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.

Running MobileFirst Platform on z13* vs x86 with similar z Systems* System of Record should increase throughput by 61% and reduce response time by 36% and front-end cost by 37%

Which platform provides the lowest TCA over 3 years?

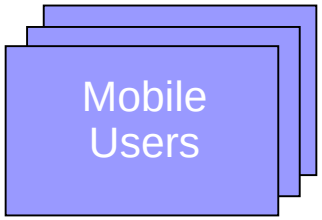

- **400 concurrent connections**
- **60% Login, 30% Add or Delete, 10% Update**

Mobile Insurance workload

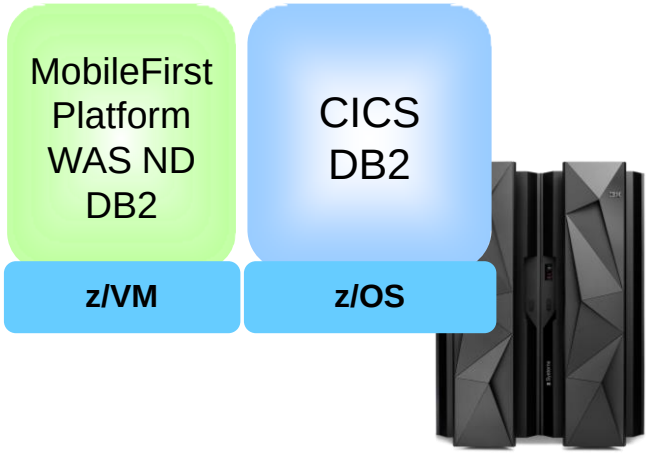

z13-409 + 5 IFLs

3,446 TPS
131ms RT
$439 per TPS
(3 yr. TCA, SoE)
Prod + Dev/QA + DR

Estimated
37%
lower cost
for systems compared

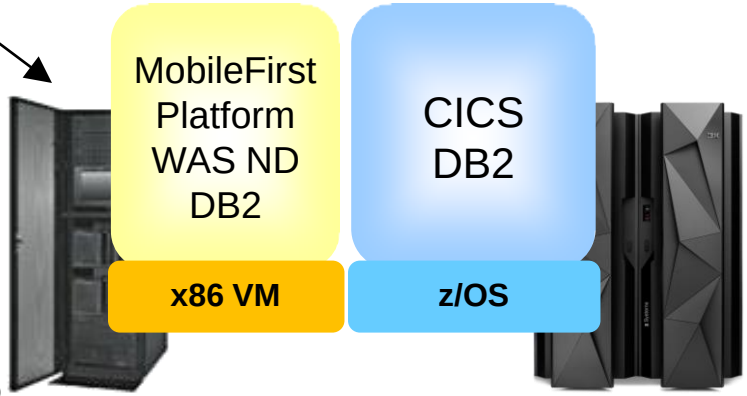

Competitor x86 System
Intel E5-2697v2 2.7GHz 6co

z13-405

2,145 TPS
205ms RT
$693 per TPS
(3 yr. TCA, SoE)
Prod + Dev/QA + DR

* 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.

Why does Linux on z perform better than Intel?

- Better hypervisor
  - z/VM overhead of 5-8% vs VMware overhead of 20-25%
- Co-location of Worklight server with the data
  - Hipersockets on z reduces response time
    For a 50 user run in both architecture we get a network IO write of 966MB/s on zLinux versus 651MB/s on Intel (so 49% faster with hipersocket) and a network IO read of 1180MB/s on zLinux versus 820MB/s on Intel (so 45% faster with hipersocket)
  - Encrypted connection between Worklight server with backend may not be required, resulting in less CPU overhead
- SMT2 enables 30% improvement in IFL throughput
- Better caching architecture in System z processors
  - We didn't measure this but expect that the multi-tiered processor cache is a performance booster

Mobile customer requirements are more demanding – fast response and data accuracy are important

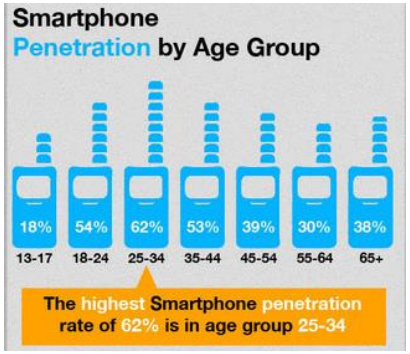

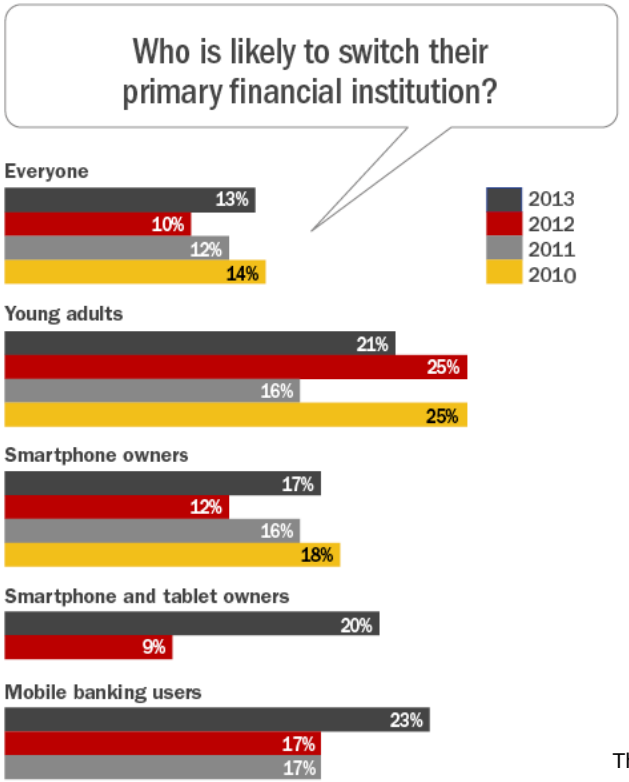

"Millenials" are potential new customers – with little patience – expect the data to be instant and up-to-date

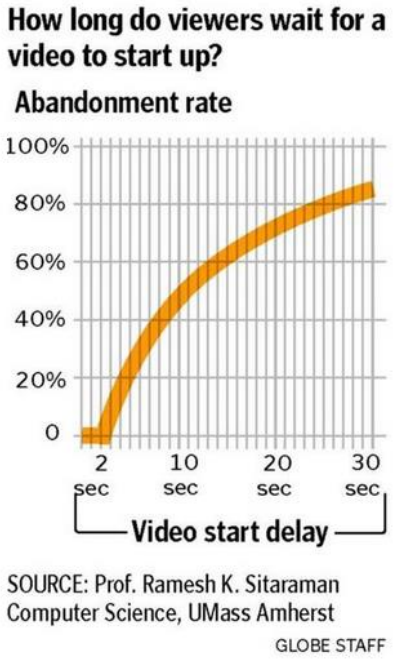

#1 reason people in the UK switch banks is dissatisfaction with the old bank's app

BBC News Sport Weather Capital Future Shop TV Radio More... Search

NEWS BUSINESS

Home | US & Canada | Latin America | UK | Africa | Asia | Europe | Mid-East | Business | Health | Sci/Environment | Tech | Entertainment | Video

Market Data | Economy | Entrepreneurship | Business of Sport | Companies | Technology of Business | Knowledge Economy

28 February 2014 Last updated at 10:03 ET

Share

Mobile banking apps recover from glitches

The mobile banking apps of several big banks have been affected by glitches.

RBS, Barclays and Santander all acknowledged problems after complaints by their customers.

RBS, which includes NatWest, said the problems with the apps were due to too many people trying to check their accounts to see how much they had been paid by their employers.

All three banks apologised but said their mobile services were now working as normal.

RBS explained via its Twitter feed that "high traffic" had been responsible for the logjam.

"We are currently experiencing record usage of our mobile banking app. Over 5,500 customers are logging on every minute," it explained.

NatWest told customers: "If you're having difficulty with our mobile app, online banking, telephone banking and ATMs are all working normally."

Santander said its mobile banking app had stalled but was now back to normal: "We apologise for any inconvenience caused while we worked as fast possible to fix the problem."

"Customers were still able to transact using Santander telephone banking, Santander branches and online banking during this time."

Barclays says use of its app is gradually returning to normal.

Last December hundred of thousands of RBS customers were left unable to use their credit and debit cards for three hours, with some saying funds had vanished from their account, because of a computer failure.

Personal Banking

Can supermarkets take on big banks?

What the new TSB means for customers

Q&A: Switching bank accounts

UK 'has lost 40% of bank branches'

Immediate access to the "System of Record" is now expected

5,500 customers logging on per minute caused slowdowns

Consumers were checking salary/benefit payments

Some had "sticky fingers"

http://www.bbc.com/news/business-26387754

IBM studied price-performance of read-only caching

- "*Pull*ing" SOR data into a mid-tier cache shields mainframe from query-only ("non-revenue producing") traffic
  - For read-only mobile workloads – mobile-initiated writes must be processed on System Of Record
  - WXS used for technology study
  - Use-case may be sub-optimal (mobile bank users want to see update data, not old)
- Tested "Sticky Finger" scenario with 500 users, each doing 20 reads per login session with 100ms think time and a 1s cache invalidation
  - Assume 500 users with 100ms think time equivalent to 5,000 users and 1s think time
- Quantified TCA savings for "Sticky Finger" scenario reading CICS/DB2 data with WebSphere eXtreme Scale (WXS) on zLinux as front-end cache
  - 64% savings for WXS caching vs business as usual (BAU)
  - 29% savings with MWP
- Quantified TCA increase for "Blended" scenario with 500 users (70% do 1 read/session, 25% do 4 reads/session and 5% do 20 reads/session)
  - Just 4% savings for WXS caching vs BAU
  - With new MWP, no caching (BAU) is 45% less cost

Architecture of read-only mobile workload with WXS caching

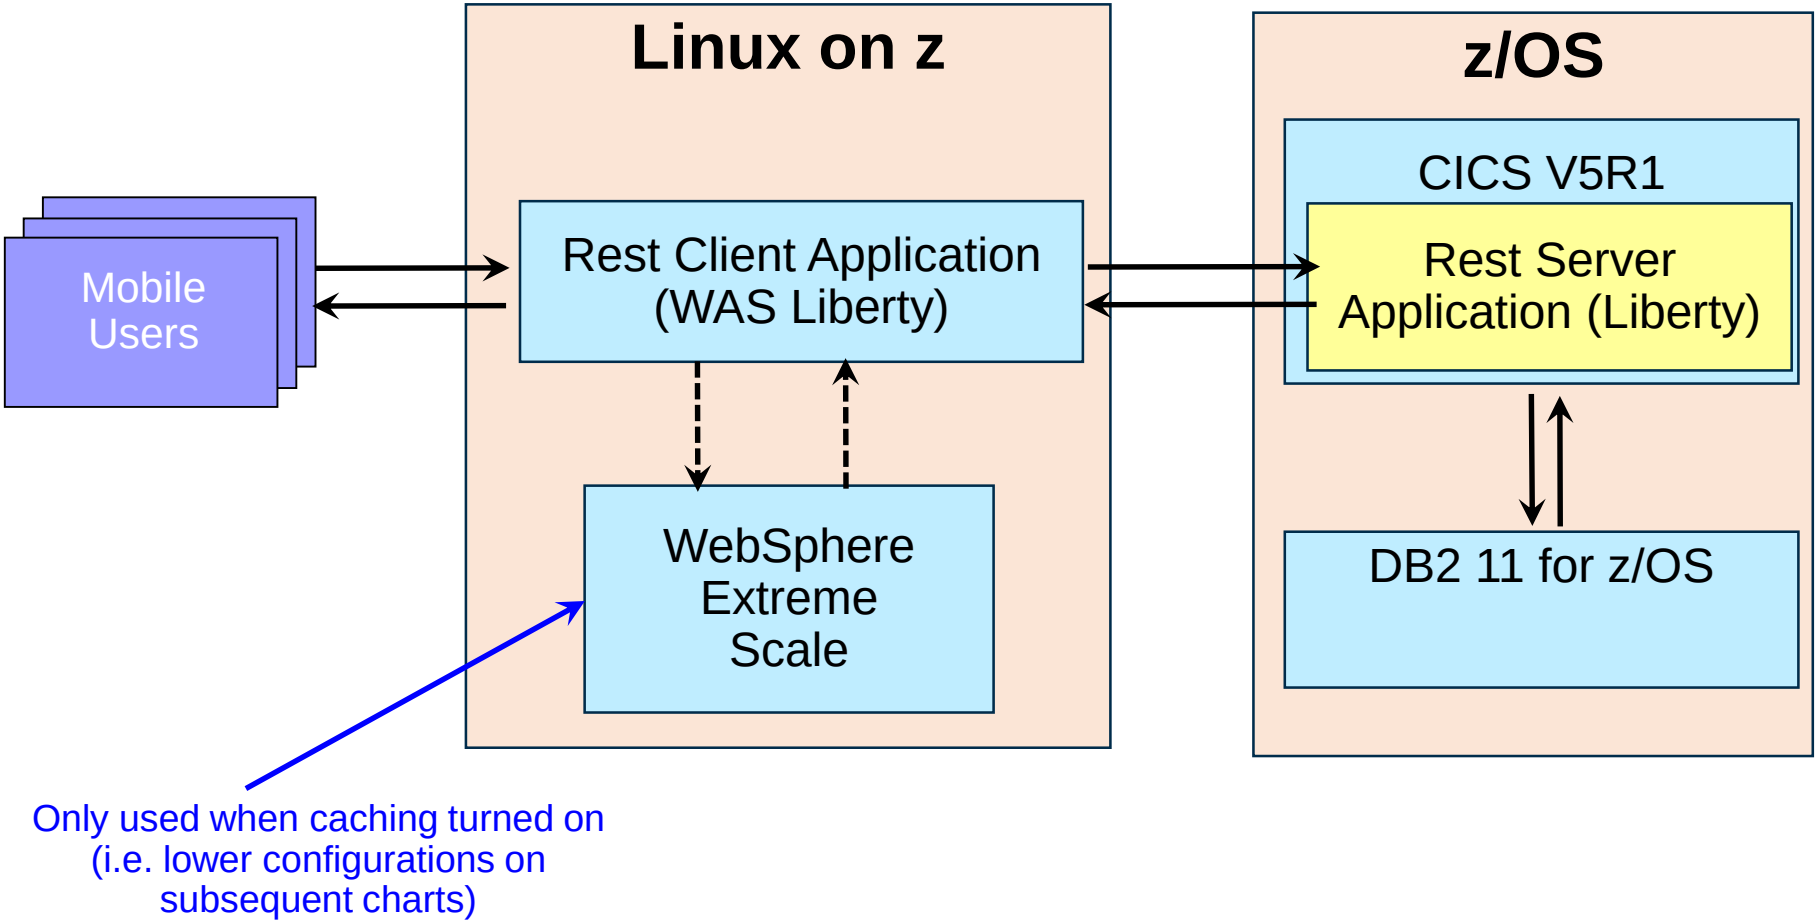

WXS caching reduces TCA for *sticky finger with think-time user* mobile workloads by 64%

Which platform provides the lowest TCA over 3 years?

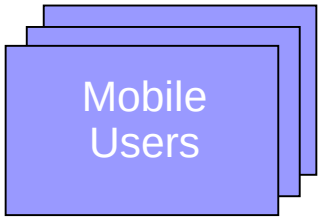

- **500 concurrent connections**
- **20 reads/session with 100ms think time (forcing cache refresh)**
- **1 second cache invalidation (WXS scenario)**

Mobile workload drives minimum throughput of **5,200** transactions per second and response time of 5ms

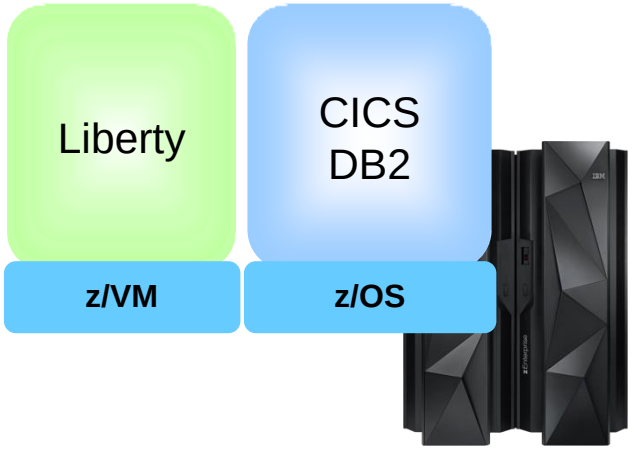

zEC12-705 + 3 IFLs

$21.8M (3 yr. TCA)
Prod
(BAU)

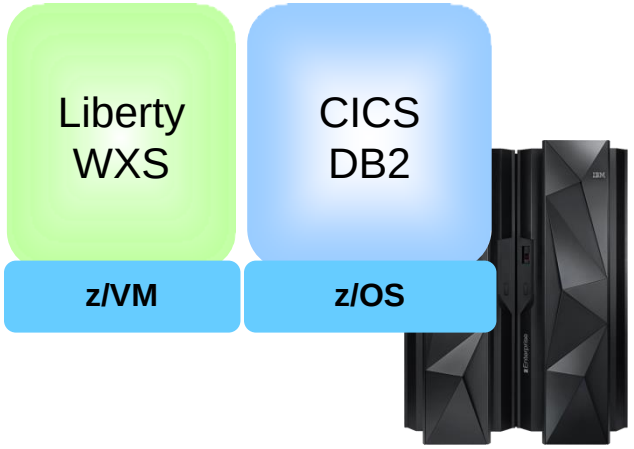

zEC12-407 + 4 IFLs

$7.9M (3 yr. TCA)
Prod

64% *lower cost!*

Mobile Workload Pricing for z/OS helps alleviate spikes

Improves the cost of growth for mobile transactions processed in System z environments such as CICS, IMS and DB2

- **IBM announced Mobile Workload Pricing for z/OS …**

  An enhancement to sub-capacity pricing

  - *Normalizes the rate of transaction growth*
  - *Mitigates the impact of Mobile on MLC charges where higher transaction volumes cause a spike in machine utilization*
  - *Works like an MSU "off-load" from a software pricing perspective, similar to Integrated Workload Pricing, not a defined price discount*

- **No infrastructure changes required (i.e. no separate LPARs) … rather an enhanced way of reporting sub-capacity MSUs**

- **Available to all enterprises running an zEC12 or zBC12 that meet the Mobile workload tracking requirements**

- **Formal announcement letter available now!**

Mobile Workload Pricing for z/OS Reporting Process

New sub-capacity reporting tool (*June 2014*)

- A new Windows-based tool is being developed to report sub-capacity MSUs and make adjustments to LPAR peaks based on Mobile transaction data
  - Standard SCRT methodology plus new feature to adjust for Mobile workload impact
  - New tool will replace SCRT for customers who take advantage of Mobile Workload Pricing

Customer requirements – *Important!*

- Tag and track Mobile transactions and produce a file showing Mobile CPU consumption each month
  - Track Mobile transactions, including CPU seconds, on an hourly basis per LPAR
  - Load the resulting data file into the reporting tool each month (IBM-specified CSV format)
  - Run the IBM-supplied tool and submit the results to IBM each month (Replaces SCRT process)

Monthly peak calculation for billing

- The tool will subtract 60% of the measured Mobile MSUs (customer input) from a given LPAR in each hour, adjusting the total LPAR MSU value for that hour
  - When an LPAR value is adjusted, all software running in the LPAR will benefit from lower MSUs
  - Tool will calculate the monthly MSU peak for a given machine using the adjusted MSU values
  - From a software pricing perspective this will look like a partial “off-load”

Provides benefit when Mobile workloads contribute to monthly peak MSUs;
Off-peak MSU adjustments will not affect MSUs used for billing by definition

Mobile Workload Pricing for z/OS helps mitigate spikes caused by increased mobile usage – *and lowers costs*

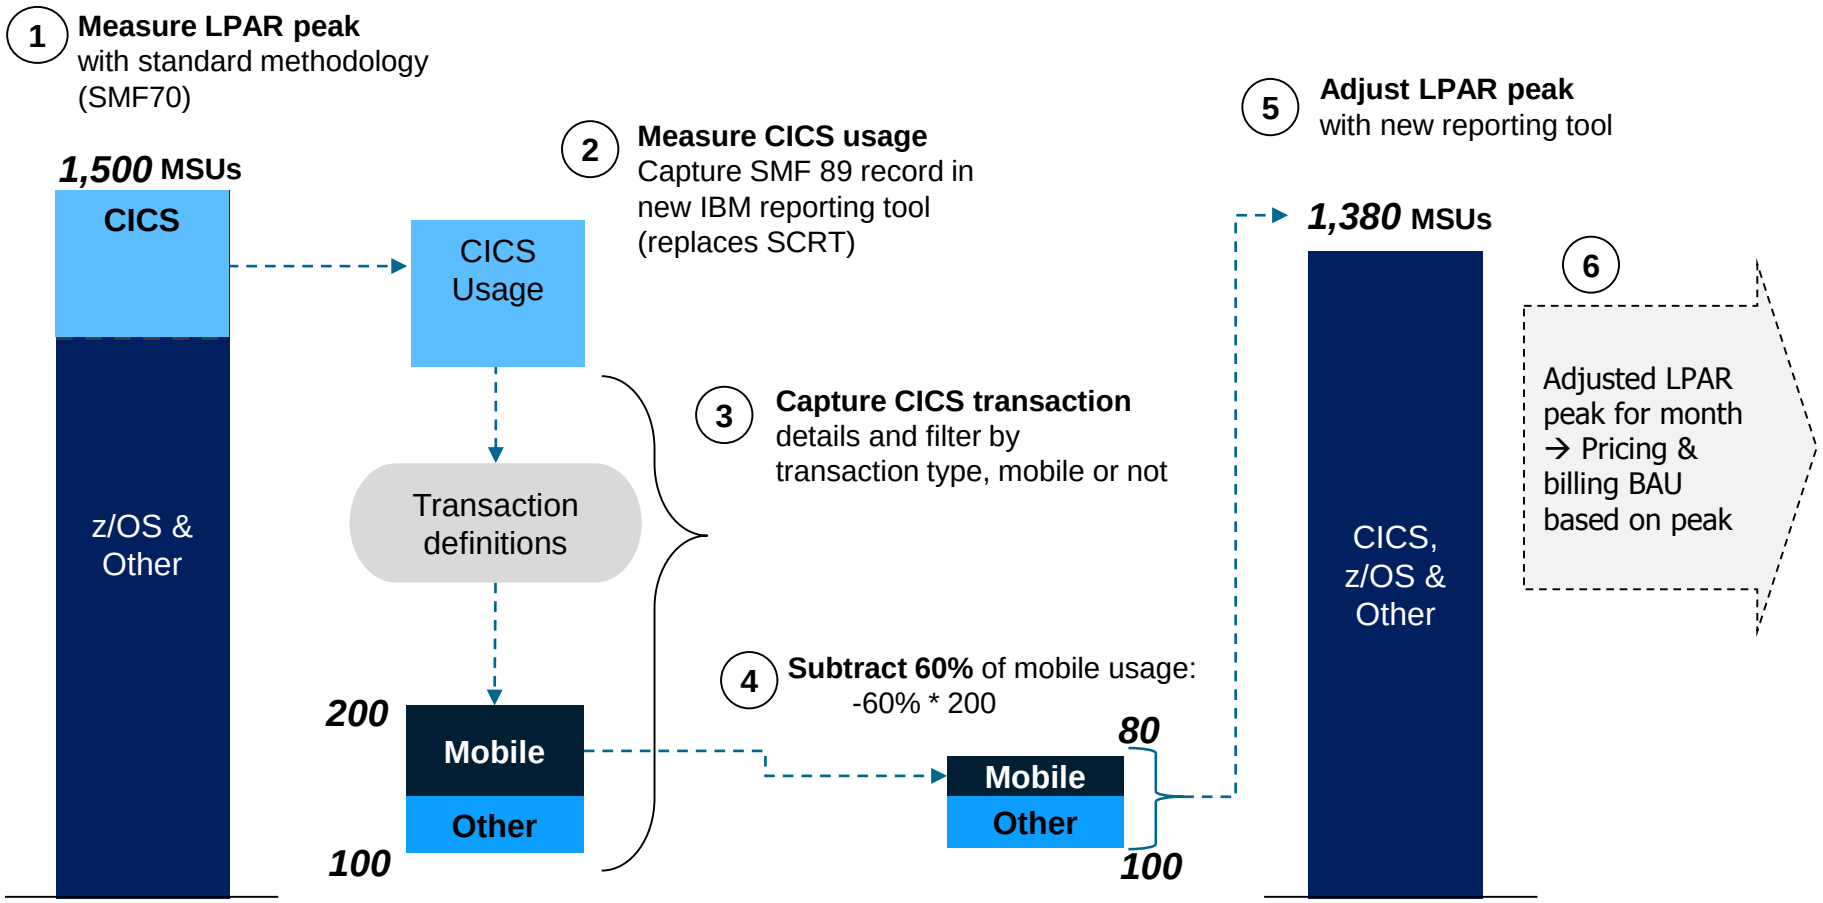

* Figures are for illustrative purposes only. Tracking process and records will vary by customer

Consider a blend of different use cases

- **70% are *regular users* (1 read/session)**
  - For these cases, the cache is ***additional overhead and cost***

- **25% are *frequent-read users* (4 reads/session)**
  - Cache shields mainframe from 3 reads
  - No cache-refresh

- **5% do 20 reads/session**
  - “Sticky finger” with cache-refresh
  - 2 mainframe reads

WXS caching reduces TCA for *blended* mobile workloads by 4% (without MWP)

Which platform provides the lowest TCA over 3 years?

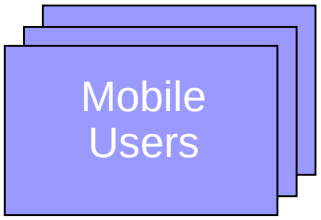

- **500 concurrent connections**
- **70% do 1 read/session; 25% do 4 reads/session; 5% do 20 reads/session with 100ms think time**
- **1 second cache invalidation (WXS scenario)**

Mobile workload drives minimum throughput of **6,300** transactions per second and response time of 12ms

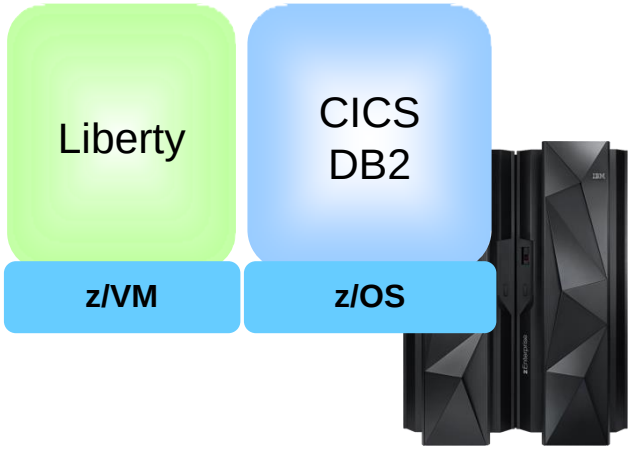

zEC12-704 + 5 IFLs
635 MSUs

$19.8M (3 yr. TCA)
Prod

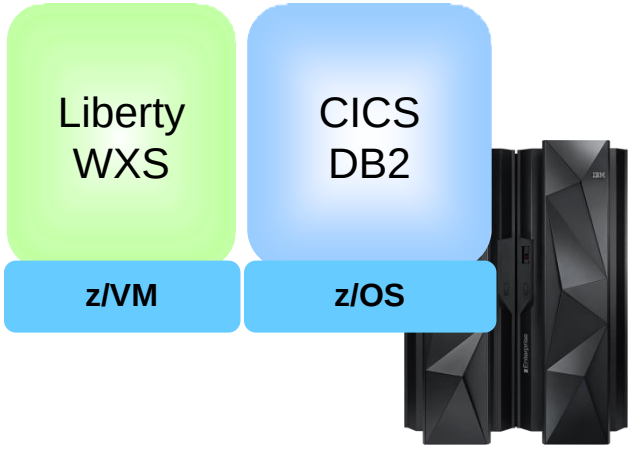

zEC12-704+ 5 IFLs
529 MSUs

$19M (3 yr. TCA)
Prod

4%
lower cost!

Mobile Workload Pricing provides 45% lower TCA than WXS Caching for *blended* mobile workloads

Which platform provides the lowest TCA over 3 years?

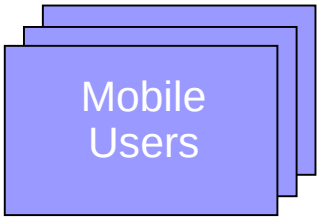

- **500 concurrent connections**
- **70% do 1 read/session; 25% do 4 reads/session; 5% do 20 reads/session with 100ms think time**
- **1 second cache invalidation (WXS scenario)**

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

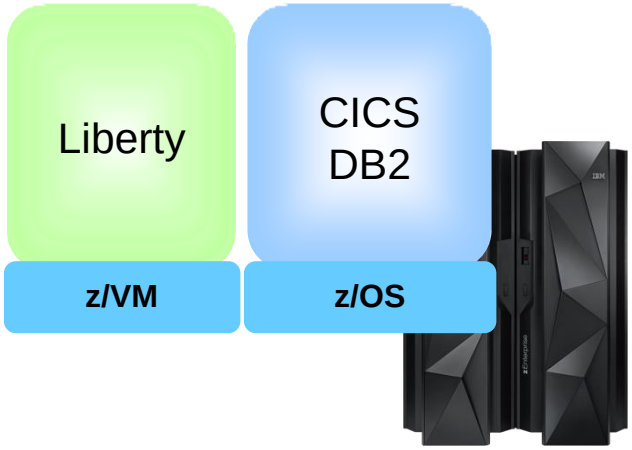

zEC12-704 + 5 IFLs
635 MSUs

$10.4M (3 yr. TCA)
Prod
Mobile Workload Pricing

45%
lower cost!

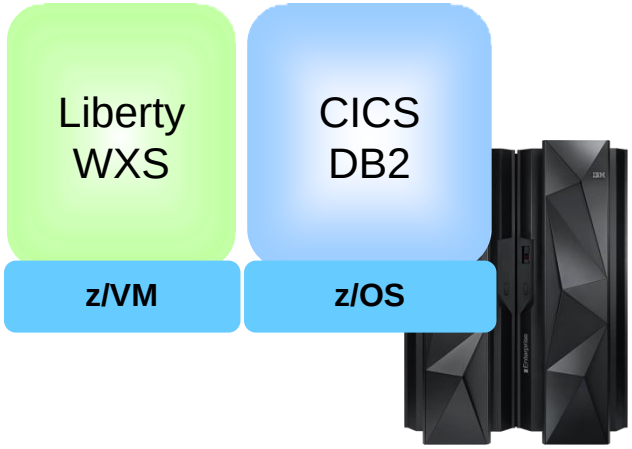

zEC12-704+ 5 IFLs
529 MSUs

$19M (3 yr. TCA)
Prod

Where best to serve SOR data to mobile devices?

1. Intermediate tier using Extract Transform Load (ETL) processing
   – Daily refresh
2. Intermediate tier using ETL with “trickle feed”
   – Intra-day updates
3. Intermediate tier using on-demand caching (IBM studied this)
   – Read-only cache (no writing BACK to the mainframe)
   – Oracle Coherence, WebSphere eXtremeScale …
4. Feed directly from the mainframe source SOR data
   – Leverage new Mobile Workload Pricing
5. Use “Push” technology
   – See Cuomo/Vila IBM Redbooks Point-of-View: Mobile Design Patterns: Push, Don't Pull

- Evaluation criteria:
  – Provide enhanced consumer experience (data freshness, responsiveness)
  – Deploy with enterprise QoS (RAS, security, scalability) at lowest cost
- Also see Williams/van der Wal Redbooks Point-of-View: System z in a Mobile World

Actual studies confirm ETL data has significant unintended costs (and on average, 12 hour-old data)

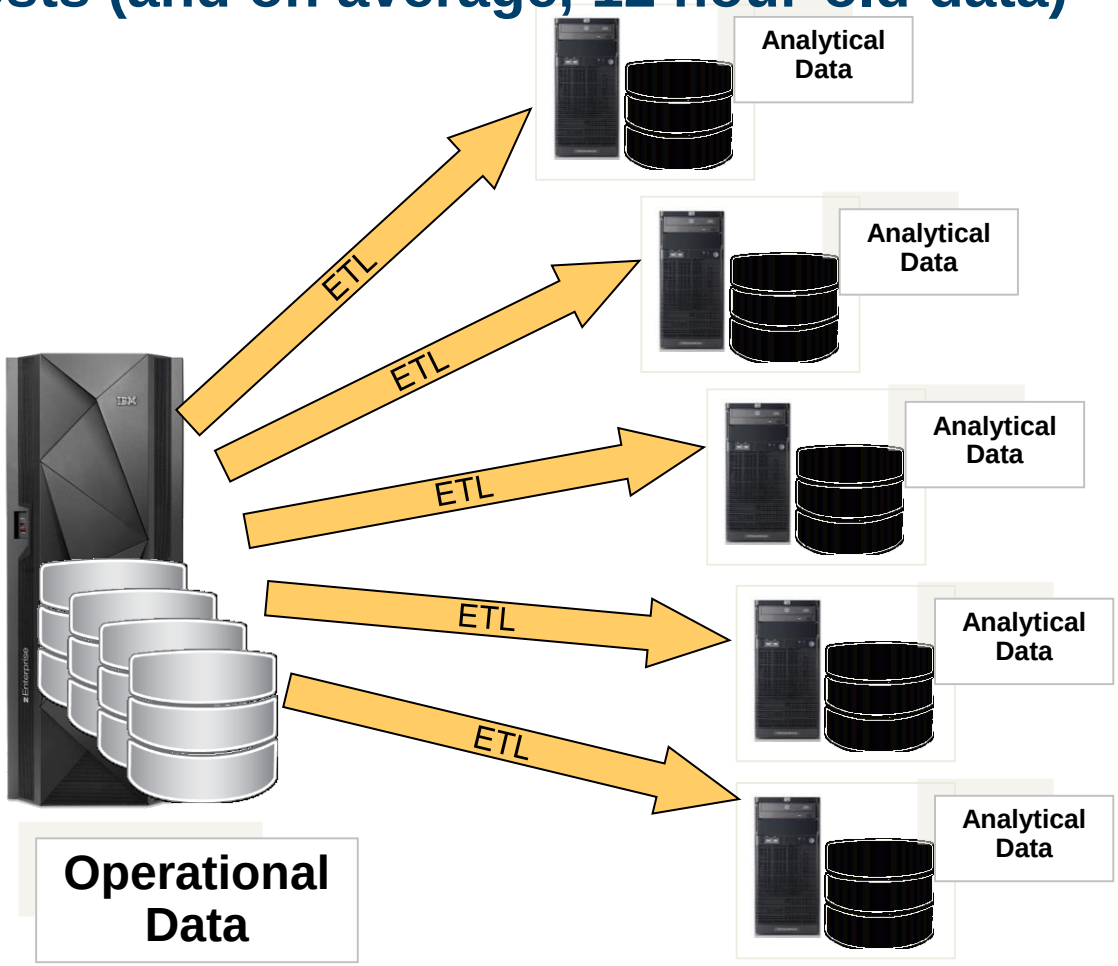

A large European bank:

- 120 database images created from bulk data transfers
- 1,000 applications on 750 cores with 14,000 software titles
- ETL consuming 28% of total distributed cores and **16% of total MIPS**

A large Asian bank:

- One mainframe devoted exclusively to bulk data transfers
- ETL consuming 8% of total distributed core and **18% of total MIPS**

ETL rule of thumb: only 20% of the data is ever used

Significant costs (often hidden) are involved when moving data off the mainframe

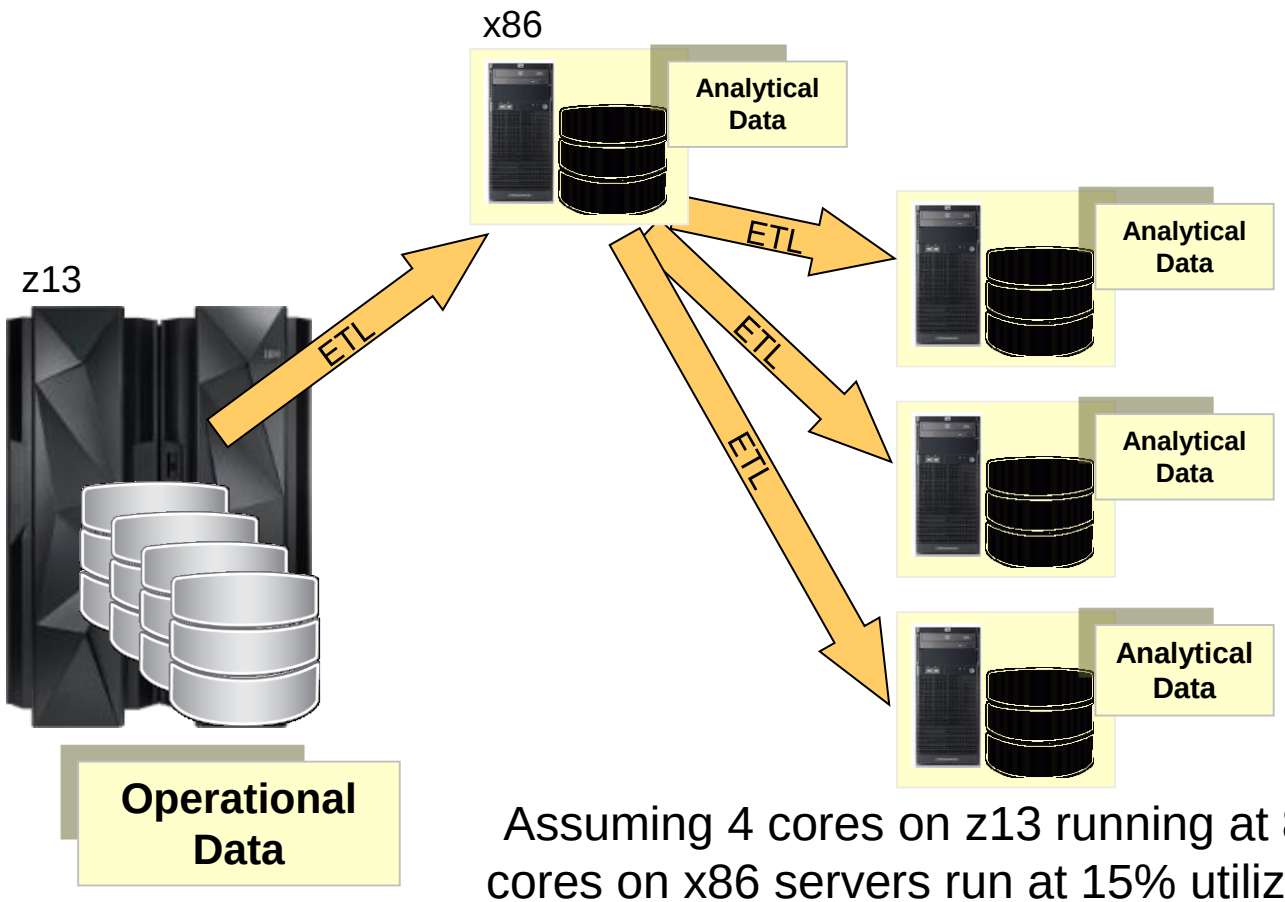

Estimated 4 yr. cost summary

System costs = $9.86M

Labor costs = $0.39M

Total = $10.3M

Assuming 4 cores on z13 running at 85% utilization and 12 cores on x86 servers run at 15% utilization, transfer will burn **519 MIPS** and use **30 x86 cores** *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 130GB file from z13 to an x86 server running with 12 x Xeon 2.4GHz E5-2440 processors. 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.

Replicated copy is unacceptable – the data is always old – resulting in excessive complexity and cost with negative ROI

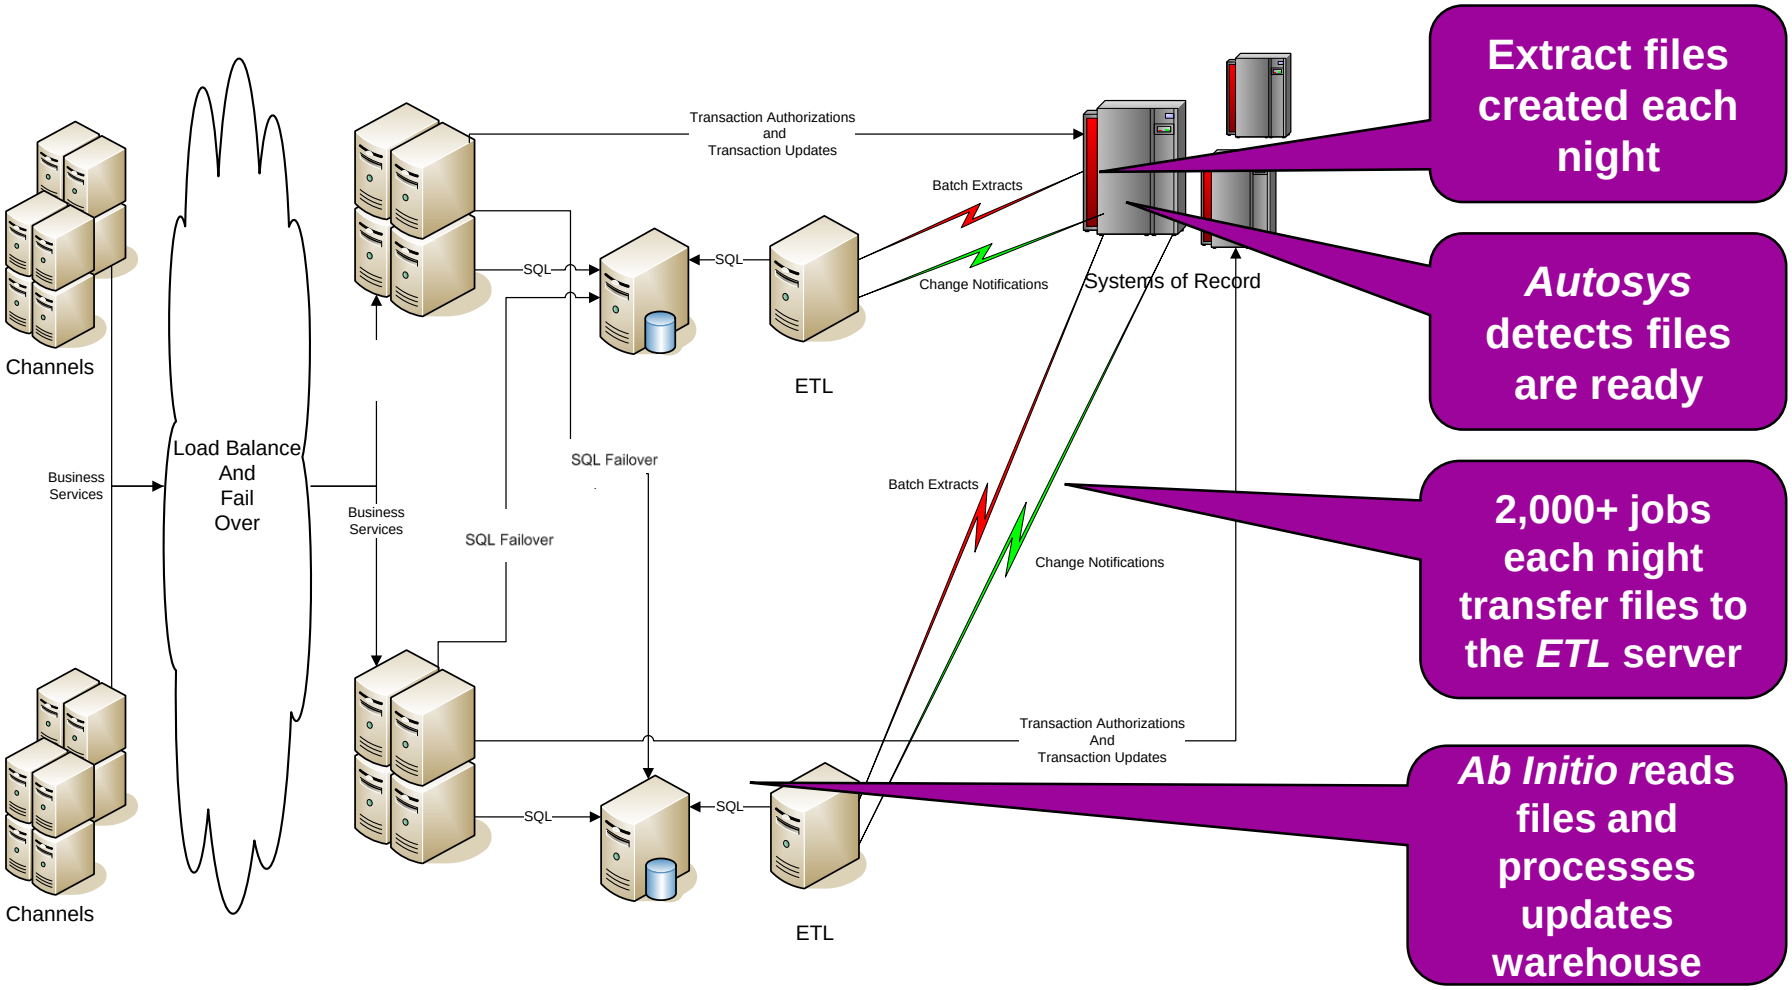

Will you reduce cost by moving your Mobile data off System z ? IBM studied* intermediate tier using ETL with “trickle feed”

Add MobileFirst Platform to create a new mobile channel to existing data and transactions.

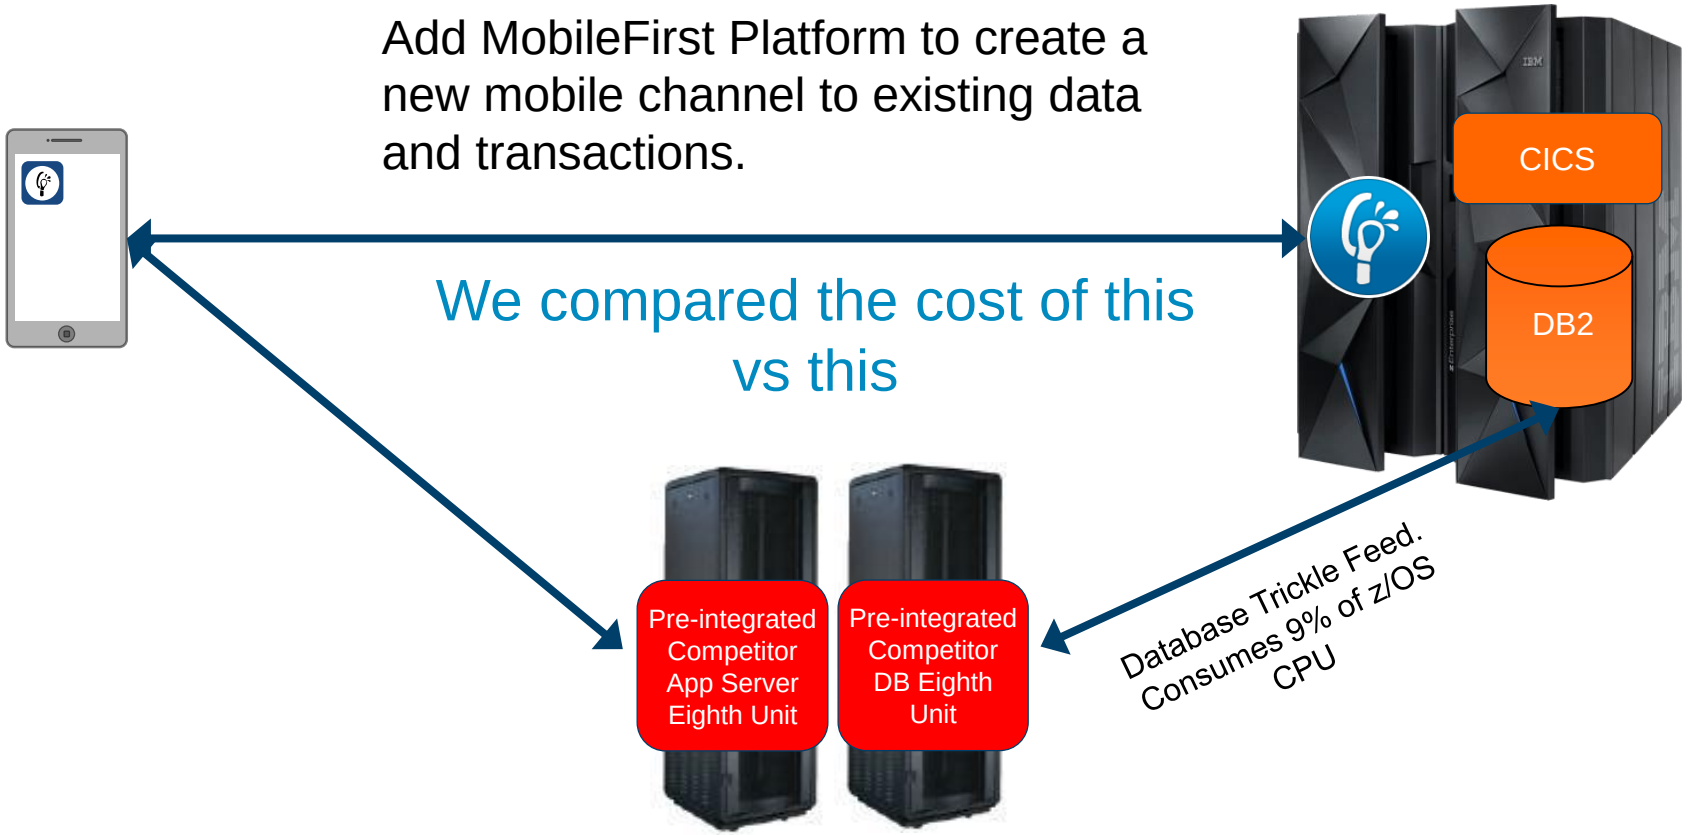

Move SoR data from System z to a new SoR off-platform, and serve mobile transactions from there.

* 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 published US list 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.

Replicating z Systems System of Record for Mobile Workloads To Competitive System is expected to increase TCA by 66% versus co-locating MobileFirst Platform and using Mobile Workload Pricing

Which platform provides the lowest TCA over 3 years?

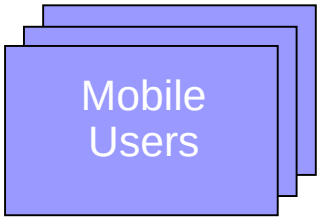

- **500 concurrent connections**
- **70% do 1 read/session; 25% do 4 reads/session; 5% do 20 reads/session with 100ms think time**
- **1 second cache invalidation**

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

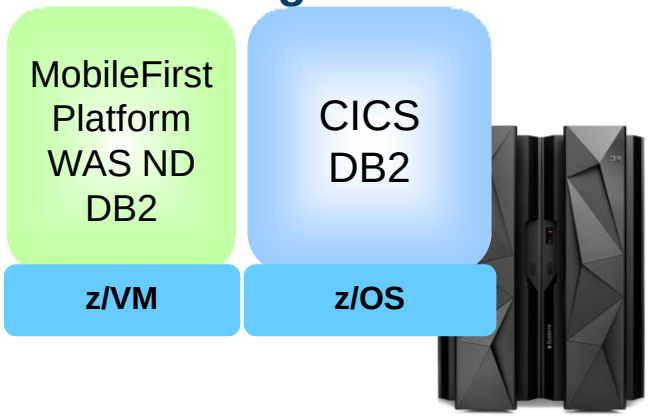

z13-606 + 4 IFLs*

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

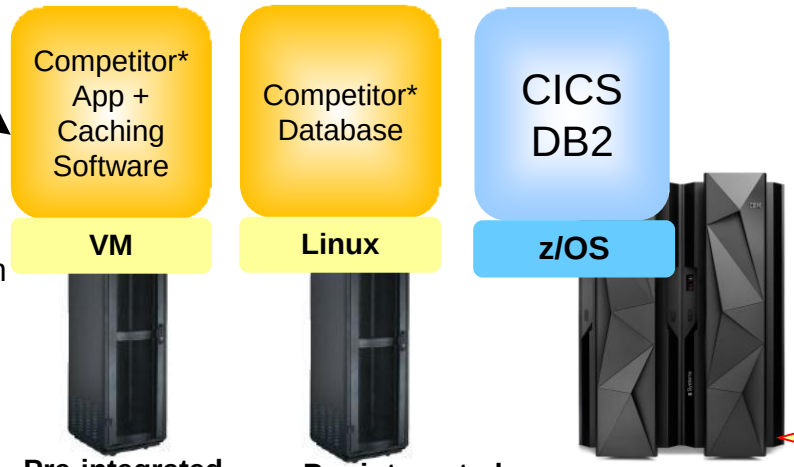

Pre-integrated Competitor* Eighth Unit (30 cores pro-rated)

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.

PUSH versus PULL study

- Cost case simulates an Automated Benefits Entitlement Deposit scenario that happens midnight end of the month

- Assumes 5% of customers are “sticky finger” users
  - 500 concurrent connections (around 5,000 TPS)

- When a deposit arrives end of the month for all 10K customers over a second of time, the throughput supported by the deposit transaction is 10,000 TPS
  - Need 2x hardware used for “sticky finger” scenario
  - To process a PUSH transaction for all 10K customers using CICS Event Processing (CEP) CPU overhead is 14%

CPU cost of CICS Event Processing

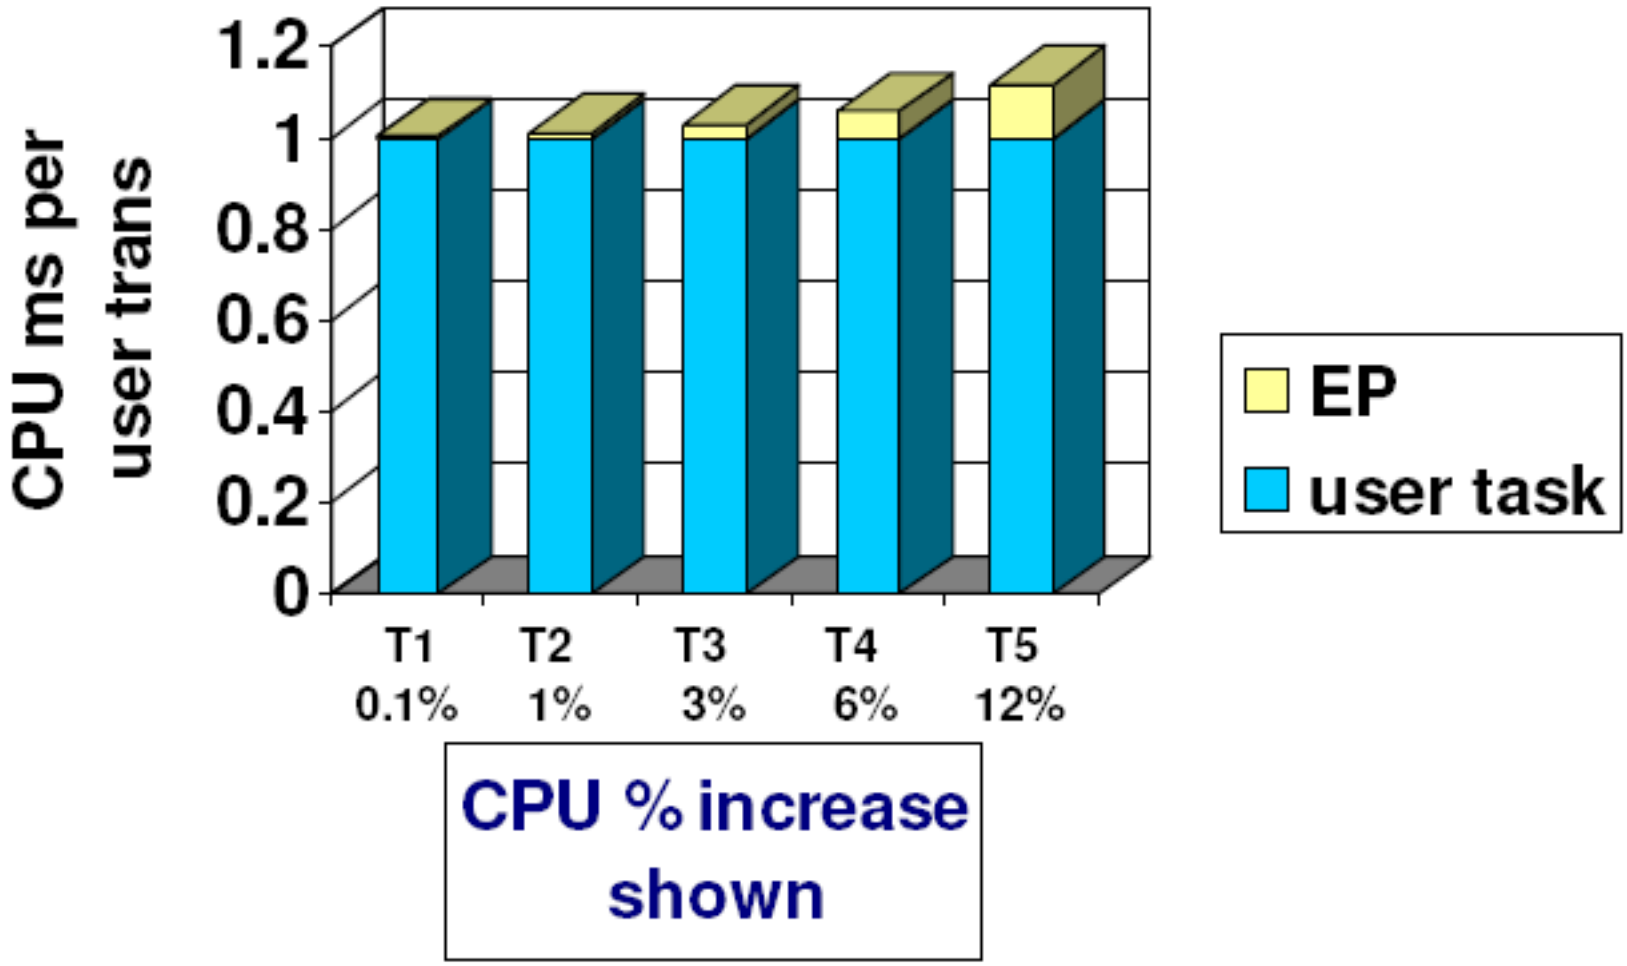

Capturing rate (per task): T1: 1 in 100, T2: 1 in 10, T3: 1 in 4, T4: 1 in 2, T5: 1 in 1

PUSH reduces TCA by 51% over PULL assuming 5% of customers with “sticky finger” behavior

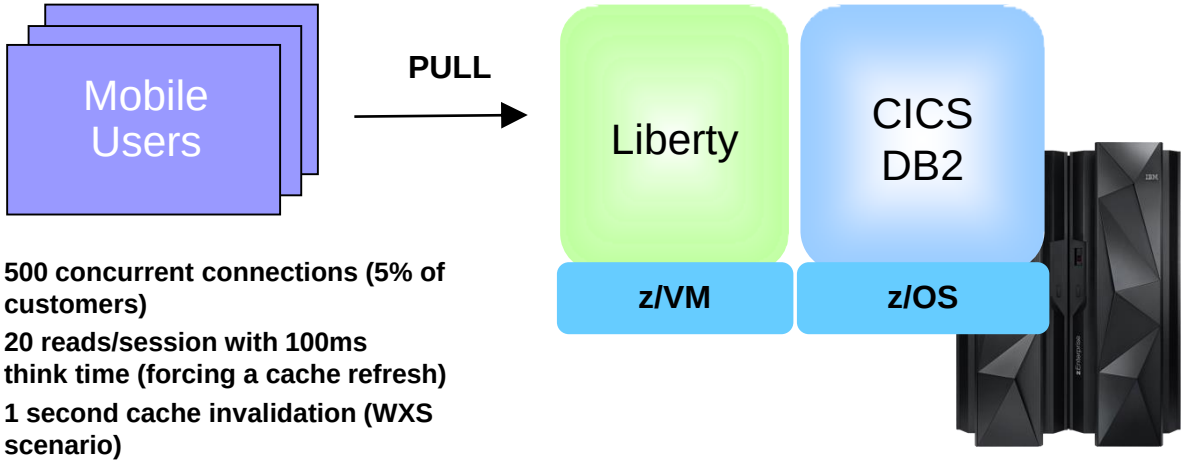

- **500 concurrent connections (5% of customers)**
- **20 reads/session with 100ms think time (forcing a cache refresh)**
- **1 second cache invalidation (WXS scenario)**

$11.2M (3 yr. TCA)
Prod
(Using Mobile Workload Pricing)

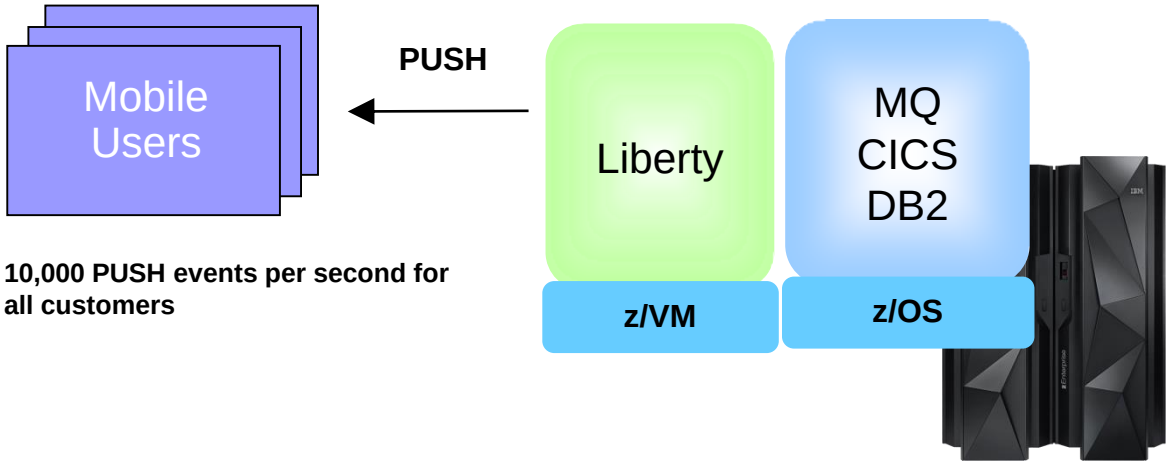

- **10,000 PUSH events per second for all customers**

$5.5M (3 yr. TCA)
Prod

51% *lower cost!*

PUSH reduces TCA by 73% over PULL for *blended* mobile workloads

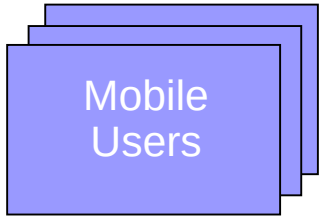

PULL

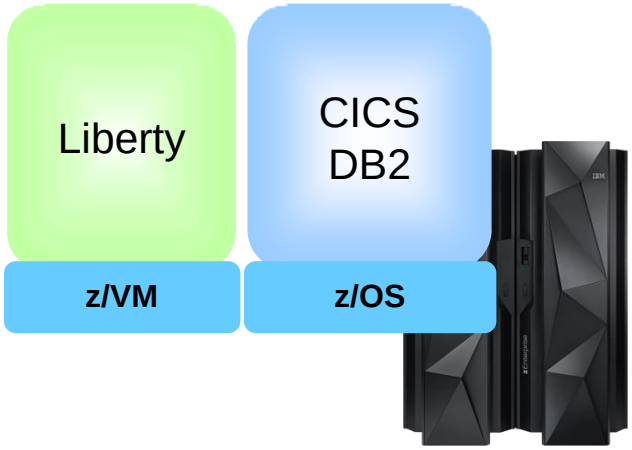

- **500 concurrent connections**
- **70% do 1 read/session; 25% do 4 reads/session; 5% do 20 reads/session with 100ms think time**
- **1 second cache invalidation (WXS scenario)**

$10.4M (3 yr. TCA)
Prod
(Using Mobile Workload Pricing)

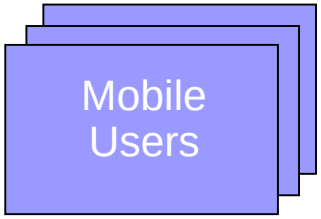

PUSH

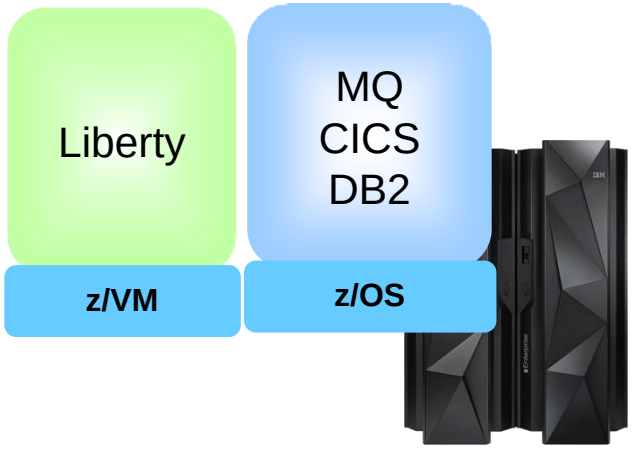

- **5,000 PUSH events per second for all customers**

$2.8M (3 yr. TCA)
Prod

73%
lower cost!

Summary

Mobile revolution is trending much more than *"eBusiness 1.0"*
- Anticipate huge volumes of transactions, with unpredictable spikes
- New opportunities, new customers

MobileFirst for Linux on z
- Extending mainframe workloads to mobile devices is ***cost-effective*** and easy
- Plus security, latency, DR advantages…

Mobile Workload Pricing
- New challenges – usage demands are tougher, less tolerant
- New mobile workloads drive monthly peaks
- Significant benefit, but initial version requires some record-keeping

Architectural choices for data serving
- **No need for an intermediate DB or cache – inefficient, costly, stale**
- *Push*ing events probably better than *pull*ing data

Questions?