



Roland Trauner

IBM Sales &
Distribution, STG
Sales

Mainframe Platform, IBM
System z Brand Manager


IBM zEnterprise EC12

Operational Analytics – Trusted Security and Resiliency – Efficiency at Scale

Betriebsanalysen

- Bewährte Sicherheit und Ausfallsicherheit

- Effizienz die skaliert





Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AIX*	FICON*	POWER7*	System z*
BladeCenter*	IBM*	PR/SM	Tivoli*
CICS*	IBM (logo)*	Smarter Banking*	WebSphere*
Cognos*	POWER*	System p*	zEnterprise*
DataPower*	Power Systems	System x*	z/OS*
DB2*	POWER4	System z10*	z/VM*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Windows Server and the Windows logo are trademarks of the Microsoft group of countries.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

* Other product and service names might be trademarks of IBM or other companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon 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 throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Das IBM zEnterprise EC12 System

Der schnellste, skalierbarste und sicherste Unternehmensserver der Welt



Parallele Geschäftsanalytik & Systembetriebsanalytik

Informationen und Einblicke
quer durch das Unternehmen

Bewährte Sicherheit und Ausfallsicherheit


Eingebaute Security Einrichtungen –
z.B. einen manipulationsresistenten
kryptografischen Koprozessor

Effizienz die skaliert

Einfache Bereitstellung von
Services in einer privaten
Unternehmens-Cloud

Die Wertschätzung von IBM zEnterprise Systemen

System z installierte Gesamtkapazität



150+

BladeCenter
Extension's
mit 1100+
Blades *

140+

Neukunden
seit der 3Q10
zEnterprise
Ankündigung

1.067

Schulen in 67
Ländern sind
Teil der IBM
Academic
Initiative für
System z

96

der Top 100
Banken
nutzen
System z

1.500

neue /
überarbeitete
Anwendungen
für z/OS und
>100 ISVs in
2011

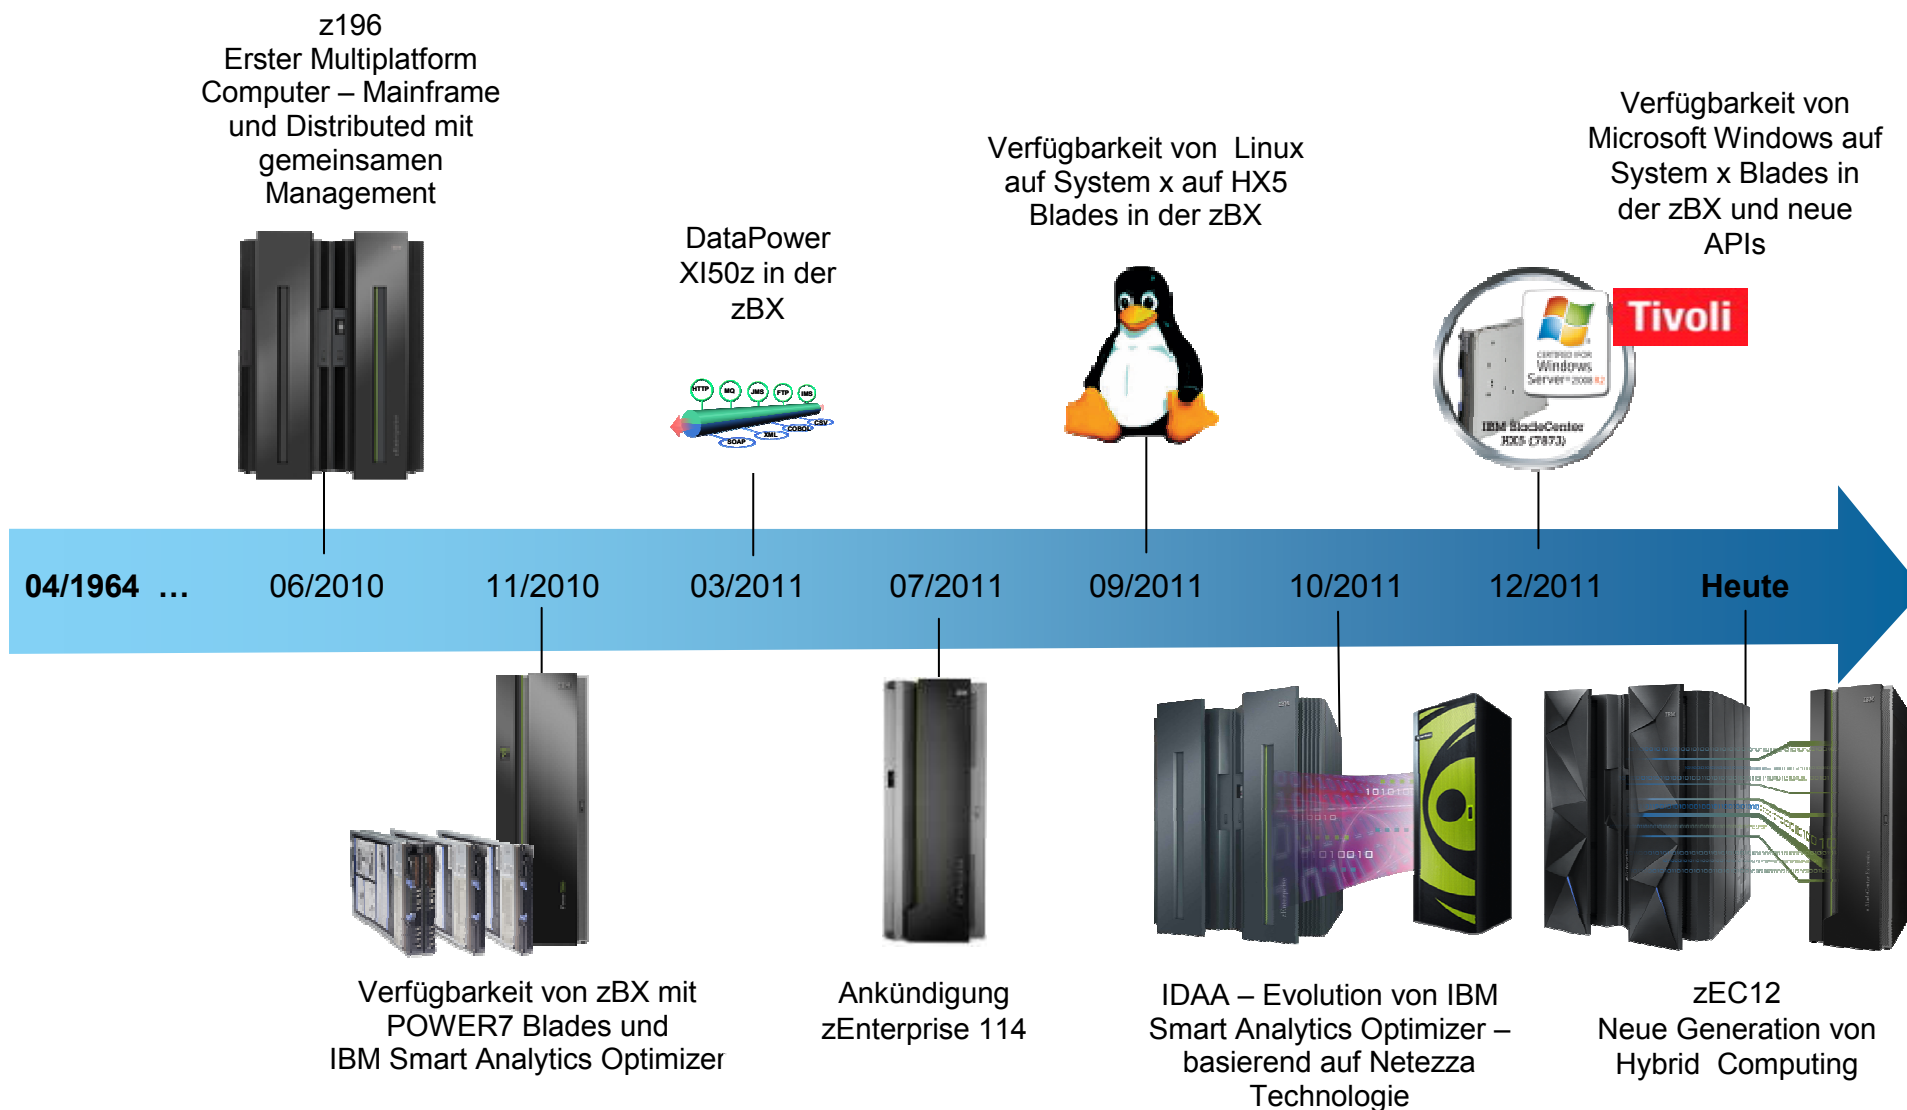
250

Kunden ww
migrierten
seit 2010
neue
Workload auf
System z

*Includes Strategic Outsourcing

**IFL = Linux-on-z Only Engine

Evolution von Hybrid Computing mit IBM System z




Ankündigung: Neue Mitglieder der zEnterprise System Familie

zEnterprise EC12 und zEnterprise BladeCenter Extension Model 003

IBM zEnterprise EC12 (zEC12)

- zEC12 nutzt den schnellsten, superskalaren kommerziellen Chip, jeder Core mit 5.5 GHz
- Innovationen zur Verfügbarkeit mit IBM zAware und Flash Express
- Optimierte für solide und leistungsstarke Datenverarbeitung in Unternehmen
- Neue Hardware Funktionen ermöglichen noch stärkere Software Performance für Java™, PL/I, DB2®



IBM zEnterprise Unified Resource Manager und zEnterprise BladeCenter® Extension (zBX) Mod 003

- Nutzt PS701 und HX5 Blades
- Liefert Ressourcen Optimierung durch Kenntnisse der Arbeitslast (workload oriented system)
- System z erweitert das Hybrid Computing

Mehr Flexibilität und Funktionen durch IDAA

- IBM DB2 Analytics Accelerator (IDAA) ermöglicht Betriebs Analysen zusammen mit dem Kernbetrieb auf der gleichen Plattform
- Analytics und OLTP kann als gleichartiger Workload betrieben werden

zBX – Eine besondere Erweiterung für die zEnterprise

- **zBX Model 003 unterstützt durch zEC12**
 - Investitionsschutz – Model 002 aufrüstbar zu Model 003
- **Keine Änderungen in den Applikationen nötig – Zertifizierungen werden von Blade-Servern übernommen**
- **Verwaltet durch den Unified Resource Manager**
 - Kundenanforderungen und Ziele bestimmen das Management verfügbarer Ressourcen
 - Ein integriertes Netzwerk hilft Sicherheit, Kontrolle und Effizienz zu erhöhen
 - Programmierschnittstellen (APIs) in Verbindung mit Systemverwaltungstools für umfassende Managementmöglichkeiten – IBM Tivoli Monitoring (ITM) wurde erweitert, um solche APIs zu unterstützen

zBX *Model: 003*



PS701 Blades

Up to
112 **HX5 Blades**

**DataPower
XI50z**

**1-4 Racks, je nach
Anzahl installierter
Blades**

Ein zBX pro zEC12

Optionale schall-
dämpfende Türen

Optionale Wärmetau-
scher für Rückwandtür

**Aufrüstbar von zBX
Model 002**


Continuing to build a multi-platform roadmap for the future

Strengthening the story

- **IBM Intends to deliver new Systems Director capability through Unified Resource Manager APIs¹**
 - Ability to discover, inventory, and visualize zBX resources under the Systems Director umbrella
 - Image management provided for blades within a zBX to support provisioning of new virtual servers
 - Energy Management of the zBX to lower energy consumption and costs
 - Power capping on System x and Power blades and power savings for Power blades
- **In the future, System z will continue to expand in the hybrid computing area¹**
- **IBM intends to deliver workload-aware optimization for System x blades through function in Unified Resource Manager¹**
- **IBM intends to deliver automated multi-site recovery for zBX hardware components based upon Geographically Dispersed Parallel Sysplex™ GDPS® technologies.¹**

¹ All statements regarding IBM future direction and intent are subject to change or withdrawal without notice and represents goals and objectives only.

zEnterprise EC12 der Kern der nächsten Generation System z



zEC12

5 Modelle: H20, H43, H66, H89, HA1

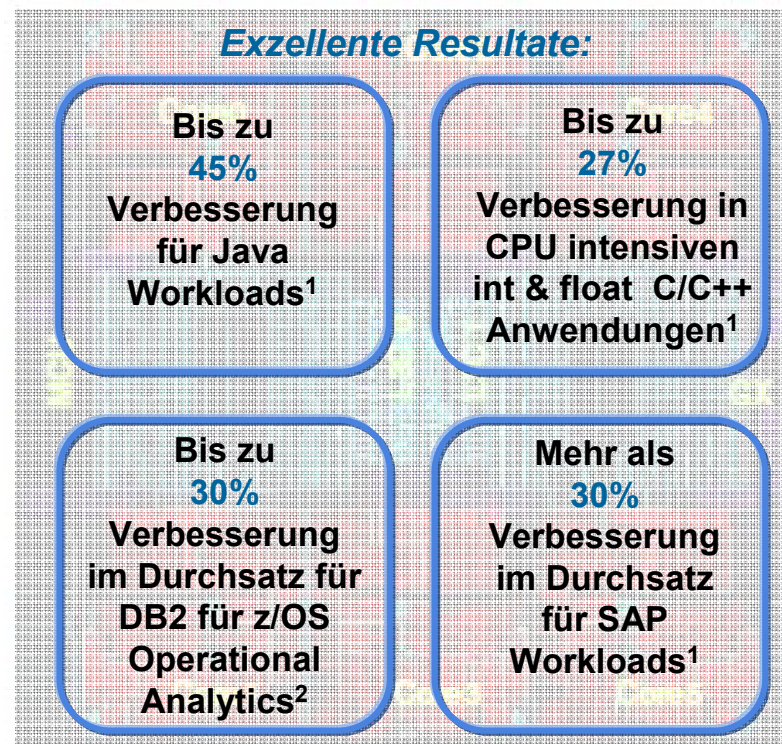
- **Prozessor Spitzen-Technologie**
5.5 GHz 6-core Prozessor Chip liefert beste Performance für alle Workloads
- **Innovation um Verfügbarkeit einen weiteren Level zu steigern**
 - **IBM zAware** bietet aktuellen Einblick in den “Gesundheitszustand” des Betriebssystems und den Anwendungen
 - **FLASH Express und “pageable large pages”** für Verfügbarkeit und Performance für kritische Workloads
- **Bewährte Sicherheit und Verfügbarkeit**
sind zEnterprise Standards

¹Basierend auf einer IBM internen Studie.

Prozessor Chip optimiert für Software Performance

Genutzt von Java, PL/I, den Compilern, DB2 und mehr

- **Ausgefeiltes Mikroprozessor Design erzeugt einen Performance-Schub für alle Workloads**
 - Zweite Generation “out of order” Design
 - Multi-level branch prediction für komplexe Workloads
- **Größere Caches um Daten Verarbeitung weiter zu optimieren**
 - Annähernd 2x auf dem Chip und 2x zusätzlich auf dem Buch
- **Neue Hardware Funktionen optimiert für Software Performance**
 - **Transactional Execution Facility** für Parallelität und Skalierung
 - **Runtime Instrumentation Facility** hilft Java Overhead zu vermindern
 - **2 GB Page Frames** bieten Performance Verbesserungen für DB2 buffer pools und Java heaps
 - Bis zu **30% Steigerung des IMS Durchsatzes** durch schnellere CPU und Cache, Compiler, und anderen¹
 - Neuer IBM Enterprise PL/I Compiler geplant um die Hardware optimal zu nutzen und um Performance Verbesserungen durch das **decimal format conversions facility** zu bekommen.



¹ Based on preliminary internal measurements and projections

² As measured by the IBM 9700 Solution Integration Center. The measured operational BI workload consists of 56 concurrent users executing a fixed set of 160,860 Cognos reports . Compared DB2 v10 workload running on IBM's z196 w/10 processors to an zEC12 w/10 processors



WIKIPEDIA
The Free Encyclopedia

[Main page](#)
[Contents](#)
[Featured content](#)
[Current events](#)
[Random article](#)
[Donate to Wikipedia](#)

▼ [Interaction](#)
[Help](#)
[About Wikipedia](#)
[Community portal](#)
[Recent changes](#)
[Contact Wikipedia](#)

► [Toolbox](#)
► [Print/export](#)

▼ [Languages](#)
[Deutsch](#)
[한국어](#)
[中文](#)

[Article](#) [Talk](#)

[Read](#) [Edit](#) [View history](#)

[Q](#)

Transactional memory

From Wikipedia, the free encyclopedia

In computer science and engineering, **transactional memory** attempts to simplify concurrent programming by allowing a group of load and store instructions to execute in an atomic way. It is a concurrency control mechanism analogous to database transactions for controlling access to shared memory in concurrent computing.

[Contents](#) [\[show\]](#)

Motivation

[\[edit\]](#)

The motivation of transactional memory lies in the programming interface of parallel programs. The goal of a transactional memory system is to transparently support the definition of regions of code that are considered a transaction, that is, that have *atomicity*, *consistency* and *isolation* requirements. Transactional memory allows writing code like this example:

```
def transfer_money(from_, to, amount):  
    with transaction():  
        from_ -= amount  
        to += amount
```

In the code, the block defined by "transaction" has the atomicity, consistency and isolation guarantees and the underlying transactional memory implementation must assure those guarantees transparently.

Hardware versus software implementation

[\[edit\]](#)

Hardware transactional memory systems may comprise modifications in processors, cache and bus protocol to support transactions.^{[1][2][3][4][5]} Load-link/store-conditional (LL/SC) offered by many RISC processors can be viewed as the most basic transactional memory support. However, LL/SC usually operates on data that is the size of a native machine word.

Software transactional memory provides transactional memory semantics in a software runtime library or the programming language,^[6] and requires minimal hardware support (typically an atomic compare and swap operation, or equivalent).



GEAR & GADGETS / PRODUCT NEWS & REVIEWS

IBM's new transactional memory: make-or-break time for multithreaded revolution

At Hot Chips last week, IBM talked about BlueGene/Q, the processor powering ...

by Peter Bright - Aug 31 2011, 11:15pm -200

109



Traditional multithreading: locks and serialization

Transactional memory is an approach to parallel programming that has the potential to make efficient parallel programming a great deal easier than it is currently. Parallel programming is easy when a task can be broken up into many independent threads that don't share any data; each part can run on a processor core, and no coordination between cores is necessary. Things get more difficult when the different parts of the task aren't completely independent—for example, if different threads need to update a single value that they share.

The traditional solution is to use locks. Every time a thread needs to alter the shared value, it acquires the lock. No other thread can acquire the lock while one thread holds it; they just have to wait. The thread with the lock can then modify the shared value (which may require a complex computation, and hence can take a long time), and then release the lock. The release of the lock in turn allows the waiting threads to continue. This system works, but it has a number of problems in practice. If updates to the shared value occur only infrequently—and hence, it's rare for a thread to ever have to wait—the lock-based system can be very efficient. However, that efficiency tends to rapidly diminish whenever updates to the shared value are frequent: threads spend a lot of their time waiting for the lock to become available, and can't do any useful work while they're waiting.

Locks also prove quite difficult for programmers to use correctly. Though the case of a single shared value is easy to handle, real programs are rarely so simple. A program with two locks, A and B, is susceptible to a problem called deadlock. If two threads need both locks, they have a choice; they can either acquire lock A followed by lock B, or they can acquire lock B followed by lock A. As long as every thread acquires the locks in the same order, there's no problem. However, if one thread acquires lock A first, and the other acquires lock B first, then the two threads can get stuck—the first waits for lock B to become free, the second waits for lock A to become free, and neither can ever succeed. This is a deadlock.

This problem might seem easy to avoid, and indeed when a program only has two locks, it normally is—but it becomes harder to ensure that every part of the program does the right thing as the program becomes more complex. Add more locks, for other bits of shared data, and it becomes harder still.

Transactional memory: the end of locks

Transactional memory is designed to solve this kind of problem. With transactional memory, developers mark the portions of their programs that modify the shared data as being "atomic." Each atomic block is executed within a transaction: either the whole block executes, or none of it does. Within the atomic block, the program can read the shared value without locking it, perform all the computations it needs to perform, and then write the value back. At the end, it commits the transaction. The clever part happens with the commit operation: the transactional memory system checks to see if the shared data has been modified since the atomic operation was started. If it hasn't, the commit just makes the update and the thread can carry on with its work. If the shared value *has* changed, the transaction is aborted, and the work the thread did is rolled back. Typically when this happens, the program will simply retry the operation.

Transactional memory potentially offers a number of advantages over the lock-based scheme. First, it's optimistic: instead of each thread needing to acquire a lock *just in case* another thread tries to perform a concurrent operation, the threads assume that they'll succeed. It's only in the case of *actual* concurrent modifications that one thread will be forced to retry its work. Second, there's no deadlock scenario, since there are no locks. Third, the programming model is, broadly speaking, one that developers are quite familiar with; the notion of transactions and roll-back is familiar to most developers who've used relational databases, as they offer a somewhat similar set of features. Fourth, atomic blocks arguably make it a lot easier to construct large, correct programs: an atomic block with nested atomic blocks will do the right thing, but the same isn't necessarily true of lock-based programs.

(It's worth pointing out that transactional memory has a number of complexities of its own: for example, what if a transaction needs to do something that can't be rolled back, like sending data over a network or drawing on the screen? The best way to approach this kind of issue, and many others, is still an area of active research.)

The hardware advantage

Up until now, transactional memory research has mostly focused on software-based implementations. Real processors don't actually support transactional memory, so it has to be emulated in some way. Some schemes make use of virtual machines to do this—there are transactional memory modifications for the .NET and Java virtual machines, for example—others use native code, and require programmers to use special functions for accessing shared data, so that the transactional memory software can ensure the right things happen in the background. A consistent feature of all of these implementations is that they tend to be slow—sometimes very slow. Although the transactional memory makes it easier to produce bug-free programs, careful use of locks (or other multithreading techniques) can yield much greater performance.

Transactional Memory: Architectural Support for Lock-Free Data Structures

Maurice Herlihy
Digital Equipment Corporation
Cambridge Research Laboratory
Cambridge MA 02139
herlihy@crl.dec.com

J. Eliot B. Moss
Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003
moss@cs.umass.edu

Abstract

A shared data structure is *lock-free* if its operations do not require mutual exclusion. If one process is interrupted in the middle of an operation, other processes will not be prevented from operating on that object. In highly concurrent systems, lock-free data structures avoid common problems associated with conventional locking techniques, including priority inversion, convoying, and difficulty of avoiding deadlock. This paper introduces *transactional memory*, a new multiprocessor architecture intended to make lock-free synchronization as efficient (and easy to use) as conventional techniques based on mutual exclusion. Transactional memory allows programmers to define customized read-modify-write operations that apply to multiple, independently-chosen words of memory. It is implemented by straightforward extensions to any multiprocessor cache-coherence protocol. Simulation results show that transactional memory matches or outperforms the best known locking techniques for simple benchmarks, even in the absence of priority inversion, convoying, and deadlock.

1 Introduction

A shared data structure is *lock-free* if its operations do not require mutual exclusion. If one process is interrupted in the middle of an operation, other processes will not be prevented from operating on that object. Lock-free data

structures avoid common problems associated with conventional locking techniques in highly concurrent systems:

- *Priority inversion* occurs when a lower-priority process is preempted while holding a lock needed by higher-priority processes.
- *Convoying* occurs when a process holding a lock is descheduled, perhaps by exhausting its scheduling quantum, by a page fault, or by some other kind of interrupt. When such an interruption occurs, other processes capable of running may be unable to progress.
- *Deadlock* can occur if processes attempt to lock the same set of objects in different orders. Deadlock avoidance can be awkward if processes must lock multiple data objects, particularly if the set of objects is not known in advance.

A number of researchers have investigated techniques for implementing lock-free concurrent data structures using software techniques [2, 4, 19, 25, 26, 32]. Experimental evidence suggests that in the absence of inversion, convoying, or deadlock, software implementations of lock-free data structures often do not perform as well as their locking-based counterparts.

This paper introduces *transactional memory*, a new multiprocessor architecture intended to make lock-free synchronization as efficient (and easy to use) as conventional techniques based on mutual exclusion. Transactional memory allows programmers to define customized read-modify-write operations that apply to multiple, independently-chosen words of memory. It is implemented by straightforward extensions to multiprocessor cache-coherence protocols. Simulation results show that transactional memory is competitive with the best known lock-based techniques for simple benchmarks, even in the absence of priority inversion, convoys, and deadlock.

TRANSACTIONAL MEMORY

Intels Haswell kommt 2013 mit neuer Speicherverwaltung

Mit seiner übernächsten CPU-Architektur, Codename Haswell, führt Intel die neue DRAM-Technik "*Transactional Memory*" ein. Vor allem Multithreading-Anwendungen sollen davon profitieren, dass sich Speicherzugriffe weniger gegenseitig behindern.

<http://www.golem.de/news>

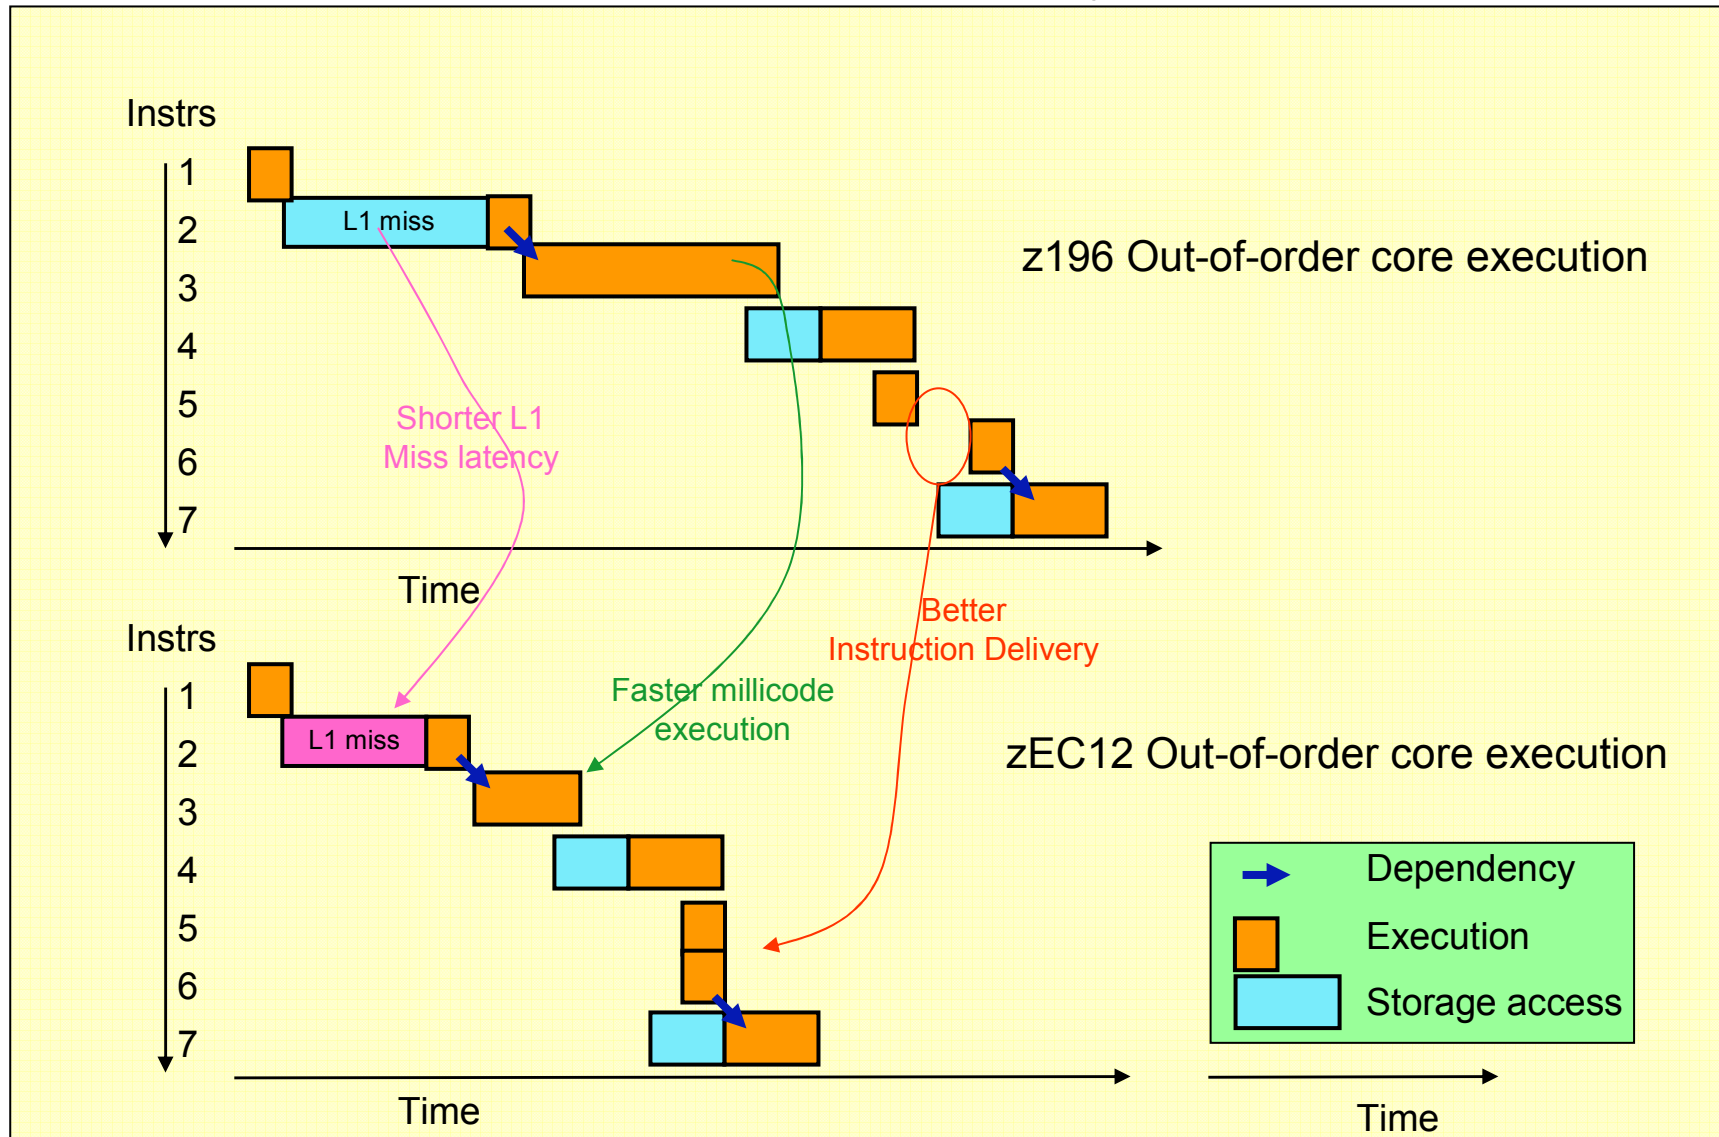
Mit Transactional Memory soll die Speicherarchitektur von [PCs](#), die sich seit Jahrzehnten kaum weiterentwickelt hat, endlich zu Mehrkernprozessoren und [Betriebssystemen](#) und Anwendungen mit vielen Threads aufschließen. Das gab Intel nun in einem [Blogeintrag](#) bekannt. Der neue Speichertyp soll mit der Architektur Haswell eingeführt werden, sie folgt 2013 der für 2012 vorgesehenen Architektur Ivy Bridge nach.



Intels Ankündigung von Haswell auf dem IDF 2011
 (Bild: Nico Ernst/Golem.de)

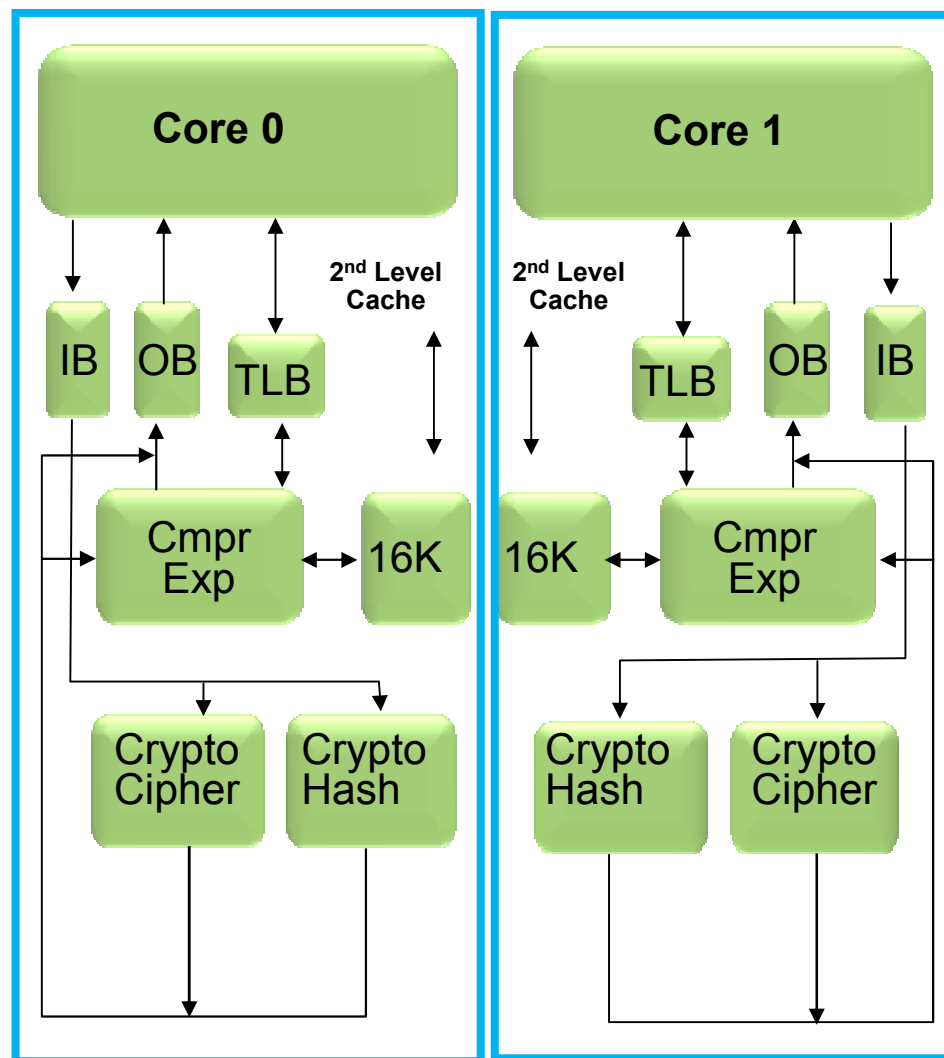
Artikel:	TRANSACTIONAL MEMORY Intels Haswell kommt 2013 mit neuer Speicherverwaltung
Inhalt:	• Drei Stufen mit neuen Befehlen
Datum:	9.2.2012, 12:04
Autor:	Nico Ernst
Themen:	Core i7 , Haswell , Ivy Bridge , Prozessor , Intel , Server , PC-Hardware
Teilen:	<div>  1  5  3 </div> <div>  Empfehlen  Twittern  +1 </div>

zEC12 OoO - Improved instruction delivery and execution




zEC12 System Compression and Cryptography Accelerator

- Coprocessor dedicated to each core (Was shared by two cores on z196)
 - Independent compression engine
 - Independent cryptographic engine
 - Available to any processor type
 - Owning processor is busy when its coprocessor is busy
- Data compression/expansion engine
 - Static dictionary compression and expansion
- CP Assist for Cryptographic Function
 - 290-960 MB/sec bulk encryption rate
 - DES (DEA, TDEA2, TDEA3)
 - SHA-1 (160 bit)
 - SHA-2 (244, 256, 384, 512 bit)
 - AES (128, 192, 256 bit)
 - CPACF FC 3863 (No Charge) is required to enable some functions and is also required to support Crypto Express4S or Crypto Express3 features




zEC12 Multi-Chip Module (MCM) Packaging

- 96mm x 96mm MCM
 - 102 Glass Ceramic layers
 - 8 chip sites
 - 7356 LGA connections
 - 27 and 30 way MCMs
 - Maximum power used by MCM is 1800W
-
- CMOS 13s chip Technology
 - PU, SC, S chips, 32nm
 - 6 PU chips/MCM – Each up to 6 active cores
 - 23.7 mm x 25.2 mm
 - 2.75 billion transistors/PU chip
 - L1 cache/PU core
 - 64 KB I-cache and 96 KB D-cache
 - L2 cache/PU core
 - 1 MB I-cache and 1 MB D-cache
 - L3 48 MB cache shared by 6 PUs per chip
 - 5.5 GHz
 - 2 Storage Control (SC) chip
 - 26.72 mm x 19.67 mm
 - 3.3 billion transistors/SC chip
 - L4 Cache 192 MB per SC chip (384 MB/Book)
 - L4 access to/from other MCMs
 - 4 SEEPROM (S) chips – 1024k each
 - 2 x active and 2 x redundant
 - Product data for MCM, chips and other engineering information



zEC12 Hexa Core PU Chip Details



- 13S 32nm SOI Technology
 - 15 layers of metal
 - 7.68 km wire
- 2.75 Billion Transistors
- Chip Area
 - 597 mm²
 - 23.7mm x 25.2mm
 - 10000+ Power pins
 - 1071 signal I/Os

- Up to Six active cores per chip
 - 5.5 GHz
 - L1 cache/ core
 - 64 KB I-cache
 - 96 KB D-cache
 - L2 cache/ core
 - 1M+1M Byte hybrid split private L2 cache
- Dedicated Co-processors (COP) per core
 - Crypto & compression accelerators
 - Includes 16KB cache
- On chip 48 MB eDRAM L3 Cache
 - Shared by all six cores
- Interface to SC chip / L4 cache
 - 44 GB/sec to each of 2 SCs (5.5 GHz)
- I/O Bus Controller (GX)
 - Interface to Host Channel Adapter (HCA)
- Memory Controller (MC)
 - Interface to controller on memory DIMMs
 - Supports RAIM design

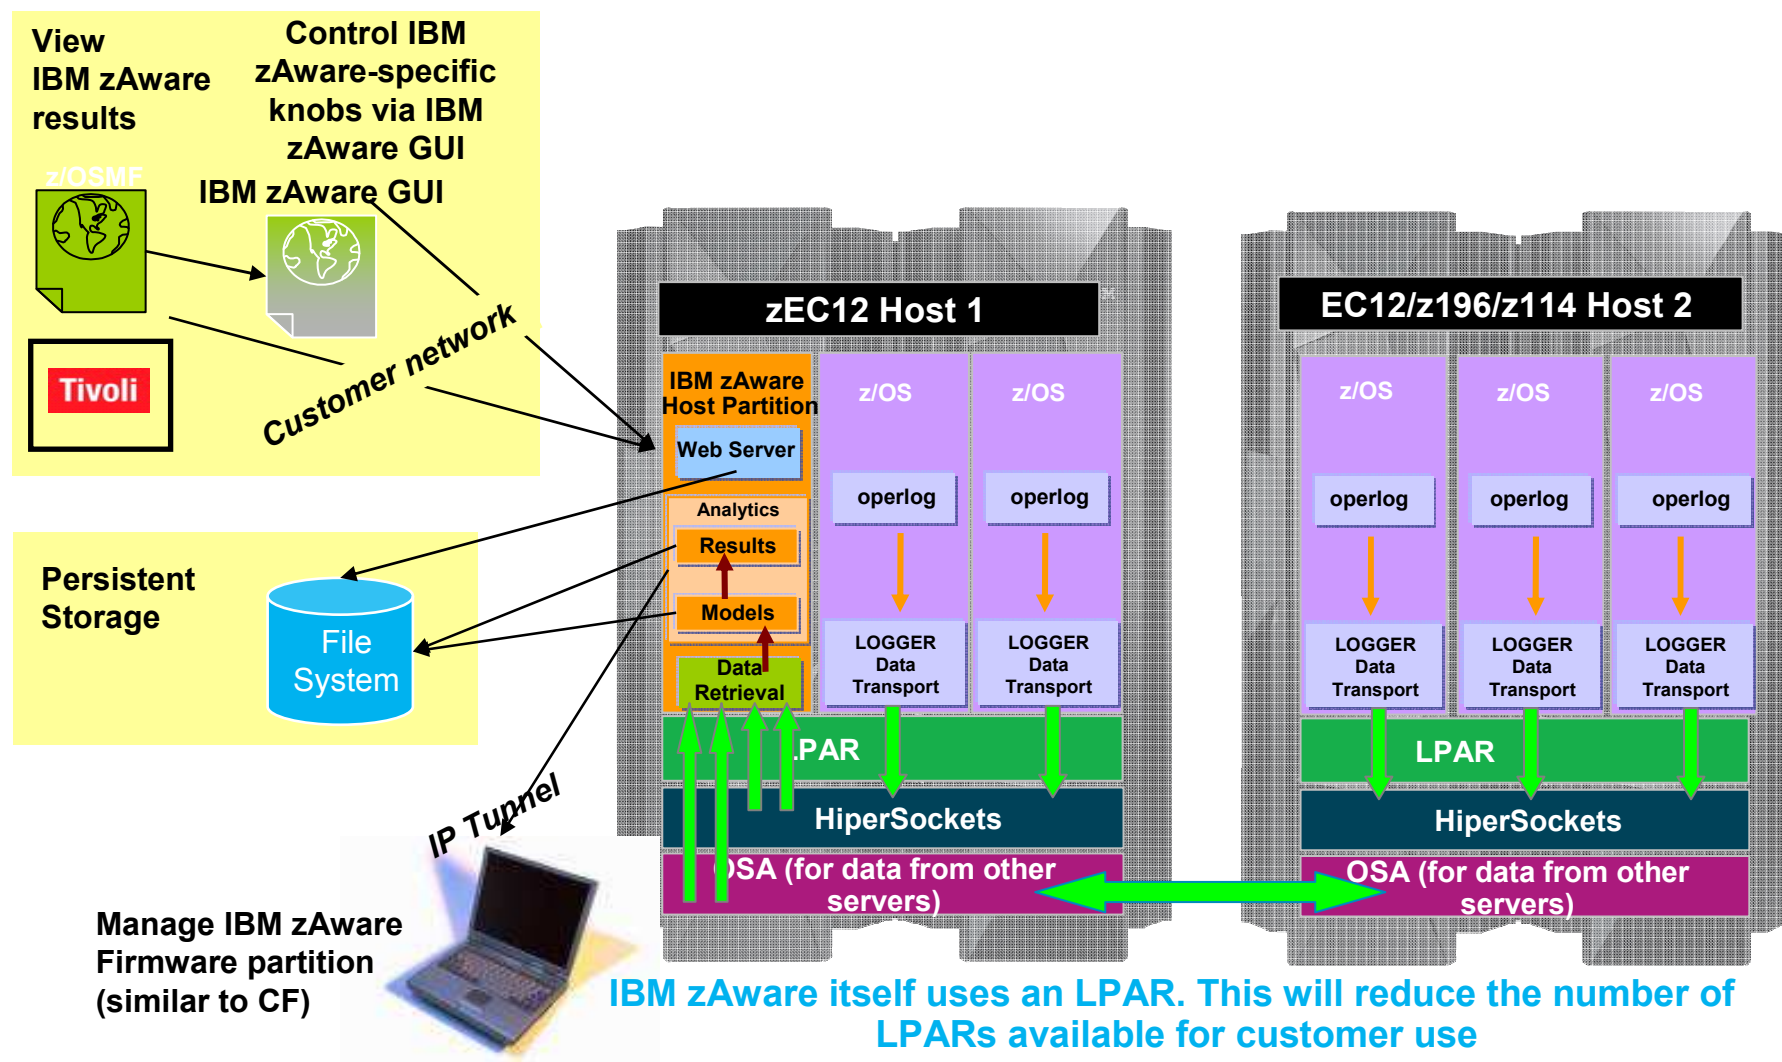
IBM zAware – Erkennt ungewöhnliches Systemverhalten

IBM zAware enthält hochentwickelte Analytik, gibt speziellen IBM insight und maschinelles Lernen, um ihr einzigartiges System besser zu verstehen.

Monitoring	Detection	Frequency	Reporting
<ul style="list-style-type: none"> ■ Unterstützung für Middleware von IBM und Fremdherstellern ■ Überwacht OPERLOG als sysplex oder monoplex ■ Fügt Auswertungen Meldungen hinzu, um Fehlerquellen frühzeitig zu erkennen 	<ul style="list-style-type: none"> ■ Erkennt Störungen, die anderen Systemen entgehen könnten ■ Erkennt seltene und ungleichmäßige Signale ■ Erkennt auffällige Mengen normaler Nachrichten ■ Erkennt kontextbezogene Unstimmigkeiten 	<ul style="list-style-type: none"> ■ Stichproben alle 2 Minuten ■ Aktualisierung der Anzeige alle 10 Minuten ■ Nutzt eine 90-tägige Datengrundlage; sammelt durchgehend neue Daten; flexibel anpassbar 	<ul style="list-style-type: none"> ■ Fast Echtzeitanalyse ■ Intuitives Reporting – Übersicht oder detailliert ■ Farbcodebasierende Browseroberfläche ■ XML Export für ISVs und andere Prozesse <ul style="list-style-type: none"> ■ Tivoli wird Alarm- und Eventbenachrichtigungen bieten




A closer look inside IBM zAware



Note: z/OS 1.13 plus PTFs or higher for monitored client

IBM Flash Express – Neue Speicherklasse

- **Flash Express ist eine innovative Lösung, entwickelt, um die Wettbewerbsfähigkeit moderner Unternehmen zu sichern**
 - Automatische Verfügbarkeitsanpassung bei kritischen Durchlaufzeiten
 - Erhöhte Verfügbarkeit und Performance für Prozesse, die keine Latenzspitzen oder wechselnde Performance ertragen
 - Latenzminimierung für kritische Anwendungen, wie z.B. Diagnose-Prozesse
- **Bietet eine sichere, belastbare und sofort einsetzbare Lösung**
- **Vorgesehene Nutzer von Flash Express und 2GB Pages:**
 - z/OS, Java, DB2, CICS, IMS
- **Robustes Design**
 - RAID10*) gespiegeltes Paar
 - Entwickelt für Langlebigkeit und “concurrent firmware upgrade”
 - Sicher : Flash Express Adapter geschützt mit 128-bit AES Verschlüsselung.



*) mirrored striping Array



Vorteile von Flash Express

- **Achieve outstanding availability**
 - Flash memory is designed to improve availability and paging performance
 - Enough capacity so that all paging data can easily reside on a pair of Flash Express cards
 - CPU performance benefits expected from the use of pageable large 1 MB pages
- **Expected Benefits**
 - **DB2 and Java:**
 - Estimated incremental several percent system CPU benefit
- **Longer roadmap**
 - Continued optimization of 1MB pageable large pages vs. 4K pages
 - Additional exploitation expected by middleware
 - Available for ISV exploitation
 - Designed for continued expanded use of new memory tier




Pages 4kB -> 1MB x 250


1MB



4kB



Pages 1MB -> 2GB x 2000



IBM System z Sicherheit als Marktstandard

- **Innere Plattformensicherheit und Privatsphäre für Transaktionen und vertrauliche Geschäftsinformationen etablieren System z als sicheren Enterprise-Anwendungsserver und Datentresor**
 - Eingebaute Hardware-Kryptographie in jedem Allzweck CP und IFL, und in den neuen Crypto Express4S Koprozessoren
 - Sichere Datenverarbeitung durch dauernden Schutz, sowohl bei ruhenden als auch bei Daten “on the flight”
- **Einsatzmöglichkeiten für Sicherheitsbedürfnissen vieler verschiedener Branchen**
 - Starker Schwerpunkt in Sicherheits- und Verschlüsselungsfunktionen, die von Banken und Finanzunternehmen benötigt werden.
 - Unterstützung der Kreditkartenindustrie mit Lösung, die Compliance und Sicherheit auf ein neues Level setzen
 - Durch eine neue Firmware für den IBM Enterprise PKCS #11 Koprozessor und Unterstützung durch z/OS, hilft der zEC12 den Ansprüchen der Europäischen Union und Kunden im öffentlichen Bereich gerecht zu werden
- **Strategischer Fokus auf die Stärken der Betriebssystemsicherheit und der kryptographischen Fähigkeiten**
 - Qualitäten, die Unternehmen mit Cloud-basierten Architekturen benötigen
 - Breite Palette an kryptographischen Basisfunktionen, nutzbar durch Betriebssystem und Middleware zur Sicherung und Beschleunigung von Workloads
- **zEC12 unterstützt die prozessorbasierte, “protected key”-basierte Verschlüsselung – exklusiv für System z**
 - Kombiniert die Leistungen der Prozessor-basierten Verschlüsselung und dem Crypto Express Koprozessor
- **PR/SM™ ausgelegt für EAL 5+ Zertifizierung**



zEC12 Overhead Power Option



Shipped separately and installed on-site to allow for door clearance



Raised Floor : Optional
Non Raised Floor : Mandatory
Co-req: Top Exit I/O option

zEC12 – Herausragende Effizienz in Rechenzentren

- **Neue, flexible Möglichkeiten für Rechenzentren ohne doppelte Böden**
- **Innovationen zur Kontrolle des Energieverbrauchs und verbesserten Effizienz in Rechenzentren**
 - zEC12 besitzt ein neues Kühlrippen-basiertes Luftkühlungssystem für eine effizientere Kühlung und verbesserte parallele Wartung
 - Durch die Wasserkühlungsmöglichkeiten der zEC12 können Rechenzentren zusätzlich bis zu 9% Energie sparen¹
 - Einsparungen mit optionaler HV Gleichstrom Stromversorgung in einem neuen Rechenzentrum werden voraussichtlich in der Größenordnung von 7-12% der Server-Eingangsleistung sein¹
- **Höhere Kapazität bei minimalen Größenänderungen**
 - Gleiche Bodenfläche der zEC12, wie die z196 und z10 EC² ohne erwähnenswerte Gewichtserhöhungen
 - Tiefe des System erhöht sich um 64 mm bzw. 2.52 in.
- **Über 12-jährige Erfahrung in Design und Konstruktion Erdbeben-resistenter Server**

¹ Based on IBM Study

² With the exception of water cooling and overhead cabling


zEC12 Water cooling infrastructure



Rear View – A Frame

2 x Exhaust
Air Heat
Exchangers

2 x Water
Cooling
Units



Rear View – A and Z Frame

Note: This is not the Rear Door Heat Exchanger design!


zEC12 - Introducing Radiator for Air Cooled System

- Closed loop water cooling N+1 pump system replaces modular refrigeration units (MRUs) used for air cooling in z196 and z10 EC

- No connection to chilled water required
- Water added to the closed loop system during installation (New Fill and Drain Tool)

- Normal operation design:


- Heat removed by water circulating to the radiator
- Fans exhaust heat from the radiator to room air




- Backup operation design

- N+1 pump/blower failure: Cooling maintained by closed loop water system without “cycle steering” slow down. Concurrent repair.
- Water cooling system failure: Cooling maintained by backup fans as in the z196 air cooled option with MRUs. “Cycle steering” slow down if needed to maintain operation. Concurrent repair

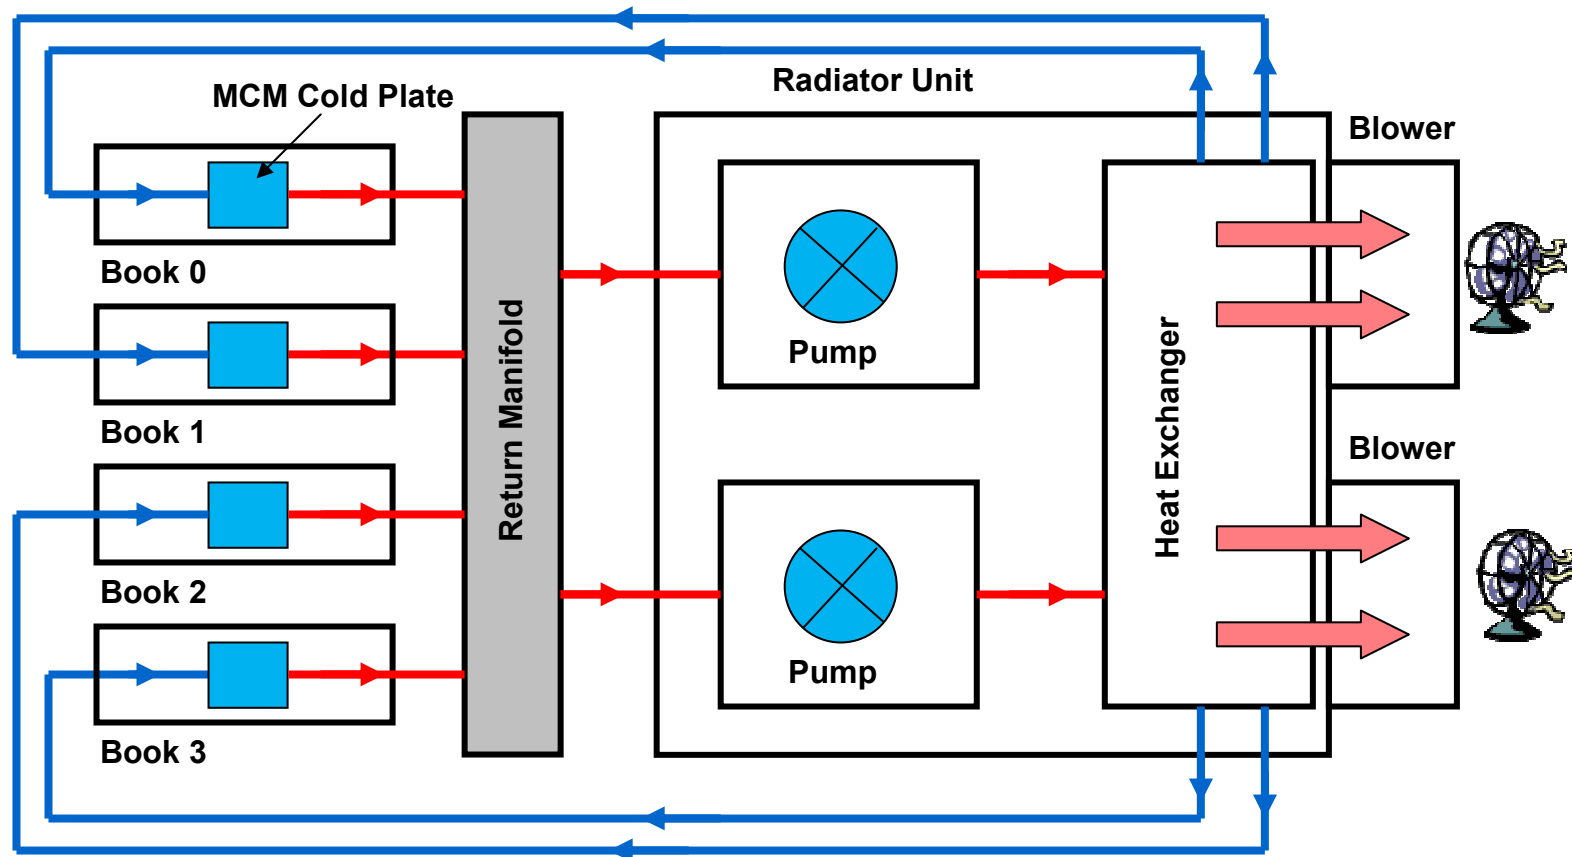
Water Conditioning Unit (WCU) detail



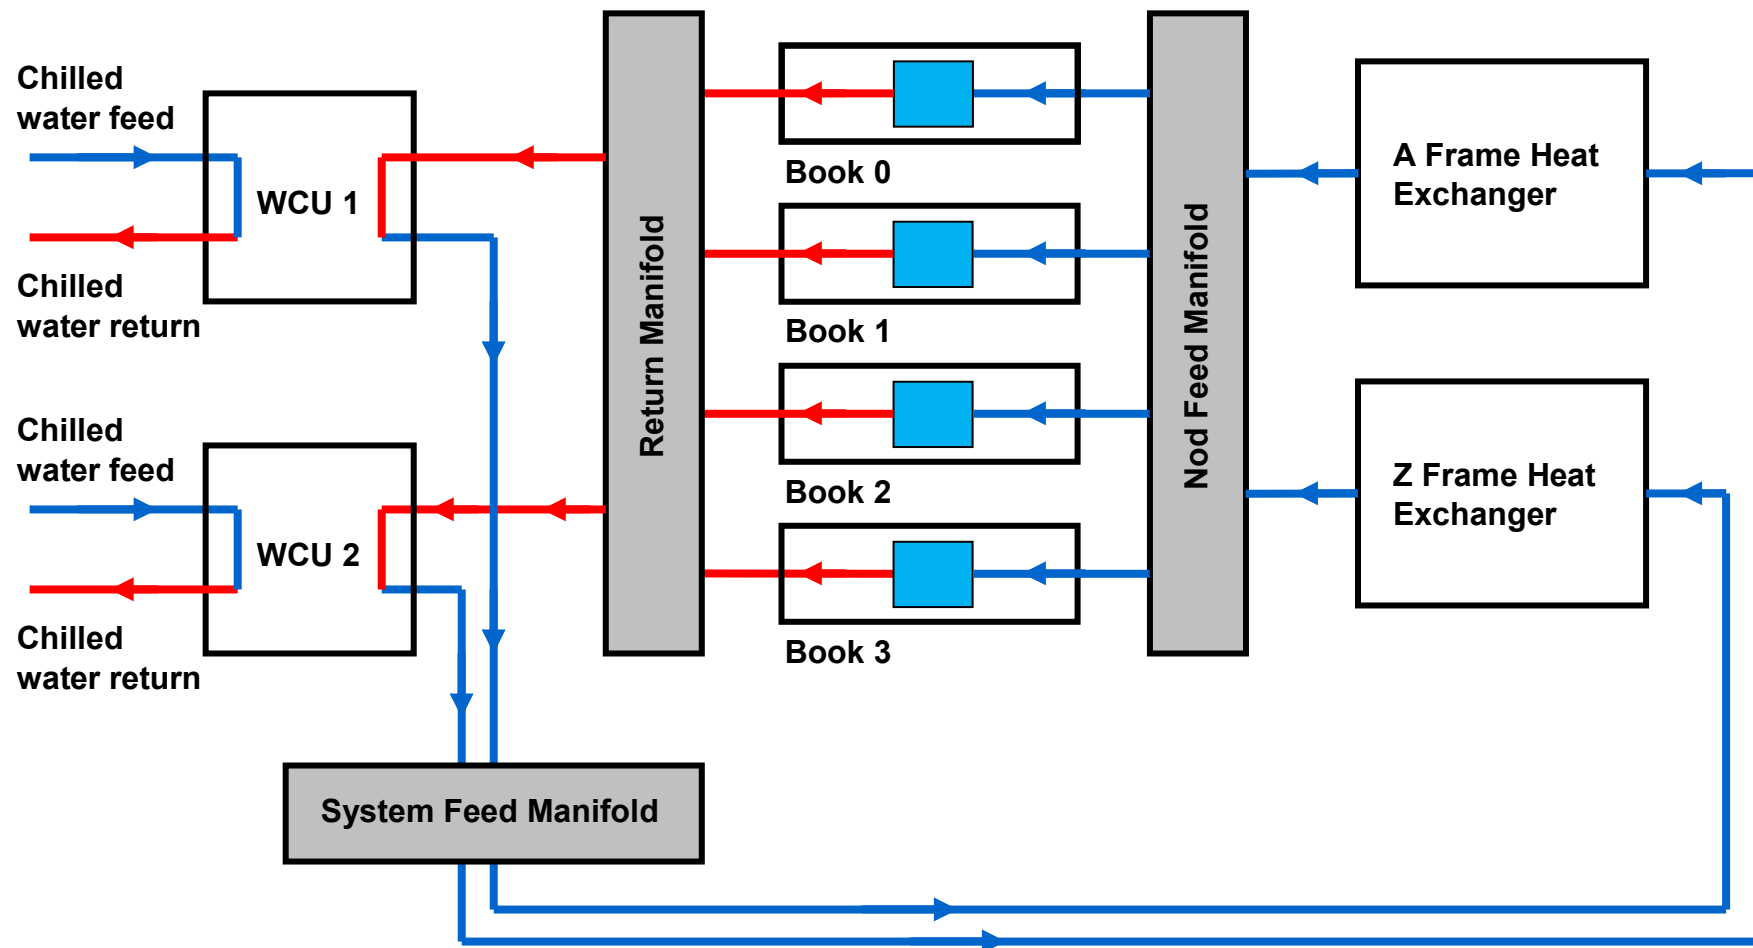
zEC12 Chilled Water Circulation - Overview



zEC12 – Radiator-based Air Cooled chilled water circulation



zEC12 Chilled Water Circulation details



zEC12 Water Cooling Requirements

- Chilled water requirements
 - Allowable system inlet water temperature range is 6-20 degrees C (43-68 degrees F), using standard building chilled water (BCW).
 - Facility chilled water quality and prevention of fouling
 - In general require only standard building chilled water without any special requirements
 - Total Hardness must not exceed 200 mg/L as calcium carbonate
 - pH must be between 7 and 9
 - Turbidity must be less than 10 NTU (Nephelometric Turbidity Unit).
 - Bacteria must be less than 1000 CFUs (Colony Forming Unit)/ml
 - Water to be as free of particulate matter as feasible
 - IBM will supply and use a deionized (DI) water solution that is mixed with benzotriazole (BTA), a corrosion inhibitor, for use within the system side cooling loop of water cooled products.
 - BTA is mixed with the Deionized water to a concentration of 1000 parts per million by weight.
 - Customer must dispose of the water solution in accordance with applicable laws and regulations and product characteristics at the time of disposal




zEC12 Customer Supplied Water Loop

The following is used for quick reference only, always use the IMPP for the latest updated information and detailed information!

- Customer Water Supply

- Designed to use standard building chilled water
- Allowable system inlet water temperature range is 6-20 degrees C (43-68 degrees F)
- Required flow rate to the frame is 3.7 – 79.4 lpm (1 -21 gpm)
- Minimum water pressure required across the IBM hose ends is 0.34 – 2.32BAR (5 – 33.7 psi)




- Supply Hoses


- IBM provides 4.2 m (14 ft) hoses, P/N 45D8967 (supply) and P/N 45D8966 (return) for use with under-floor water supply connections.
 - Customer end of the hoses are left open, allows customer ability to "cut-to-length" to match installation requirements. Insulation clamp P/N 15R9363 is provided to secure the insulation and protective sleeving once the hose has been cut to the desired length and/or installed onto customer facilities plumbing.
- IBM recommends the use of shut-off valves in front of the hoses. Allows for removal of hoses for a service procedure or relocation.
- Stainless steel fitting P/N 44V6056 available for ordering. Fitting is barbed on one side and has a 25.4 mm (1 in) male NPT. Not supplied with original ship group, must be ordered separately.

- **zEC12 water supply requirements are different from zBX Rear Door Heat Exchanger!**
- **Installation of water cooling connections to data center can result in additional installation time due to scheduling, preparation and installation of hoses, water sensors, etc.**

zEC12 Radiator Assembly for Air Cooled System




zEC12 Radiator cooling components



zEC12 Radiator Field Replaceable Units (FRU's)

- All electronics, pumps, blowers, sensors are fully redundant and concurrently replaced without cycle steering
- If plumbing manifolds or heat exchanger ever fails, air backup keeps server up



zEC12 Radiator Cooling Unit for Air Cooled System

zEC12 Radiator Cooling Units




Front View
Internal system water closed-loop



Rear View
Customer Supplied Water

Connections to the Books
(Internal system water closed-loop)



Comparing MRUs with Radiator Unit for zEC12

z196 MRU Front



zEC12 Radiator Front (NEW)



z196 MRU Rear



zEC12 Radiator Rear (NEW)



zEC12 Fill and Drain Tool (FDT)



System Fill Procedure

- Driven through Repair & Verify on SE
- 15-20 minute procedure
- Initial setup includes:
 - Starting R&V
 - Gathering FDT and BTA water solution
 - Plugging FDT into bulk power port on system

Approximate FDT unit dimensions:

- 35 inches from floor to top of handle
- 30 inches long
- 22 inches wide

Water Fill and Drain Tool – Required for both Air and Water Cooled Systems

- z196 FC3377
 - Water Cooled Models only



- zEC12 FC3378
 - Water Cooled Models
 - Air Cooled Models (radiator)



Fill and Drain kits are different on z196 and zEC12 servers and are NOT cross compatible!


zEC12 Water Cooling Considerations

Optional Water Cooling

- MES from Water Cooled (WC) to Air Cooled (AC) – Not available
- MES from Air Cooled (AC) to Water Cooled (WC) – Not available
- Requires customer supplied water, designed to use the same water supply as is currently being used by the Computer Room Air Conditioning (CRAC) units or standard building chilled water.
 - Customer water source from two different chillers are required (recommended)
 - IBM responsible for handling internal water only (filling and draining)
 - Disposal of any internal system water is Customer responsibility. This should be done in accordance with local regulations.
- Fill Drain Kit FC 3378 required, ordered via eConfig
- CPC and zBX with Rear Door Heat Exchanger, water cooling requirements are different

zEC12 Cooling

- Air Cooled System
 - Normal Air exhaust front to rear of frame
 - Fully redundant radiator cooling replaces MRUs used in z196/z10 EC
 - If complete failure of radiator system occurs, air backup keeps server running
 - In air cooling mode, hot air exits through top of system, cycle steering mode possible
- Water Cooled System
 - Normal Air exhaust front to rear of frame
 - Will operate with a single WCU, air cooled if no WCU or customer water available
 - If complete failure of water cooling system occurs, air backup keeps server running
 - In air cooling mode, hot air exits through top of system, cycle steering mode possible



zEnterprise Momentum in IBM Wachstumsinitiativen

Smarter Planet

IBM Intelligent Operations Center for Smarter Cities verbindet Datenvisualisation, Zusammenarbeit in Echtzeit und tiefgreifende Analyse, damit sich Stadtverwaltungen auf Probleme vorbereiten können, bevor sie auftreten.

Analytics

IBM Smarter Analytics: Betrugs- und Missbrauchs- bekämpfung für Versicherungen und Kliniken lernt durch neueste Daten, operationale Effizienz zu optimieren und Rückerstattungen von widerrechtlichen Schadensmeldungen zu minimieren.

Cloud

IBM Server Implementation Services for Cloud – Linux on z/VM entwickelt zur Planungs- Design- und Implementierungs-Unterstützung von Anwendern für eine flexible und robuste, Linux-basierte, private Cloud mit gutem ROI

Growth Markets

zEnterprise EC12 Design bringt Technologien, die besonders für Unternehmen in Wachstumsmärkten attraktiv sind: Keine benötigten Doppelböden und höhere Leistungsfähigkeit bei gleichem Energieverbrauch verglichen mit z196.

Deploy IBM zAware for high availability with confidence



Product New Features:


- GTS service designed to assist clients to assess, design and implement IBM zAware in a System z environment to provide protection against rare anomalies, sick-but-not-dead events and inadvertent mistakes. We also help to integrate IBM zAware with client's existing high availability and monitoring environments.

Features / Business Value:

- Provide assessment, planning, education and implementation of IBM zAware.
- Deliver processes, tools, skills and best practices that enable a fast and accurate implementation.

Services Benefits:

- Increase ROI with smooth and faster deployment of IBM zAware leveraging repeatable methodology, automation tools, and project management.
- Assess client's existing infrastructure and integrate IBM zAware seamlessly into client's existing high availability or monitoring structure with expert skills.



Learn More: <http://www-935.ibm.com/services/us/en/it-services/server-product-services-for-system-z.html>

IBM Server Implementation Services für Cloud – Linux unter z/VM

Help accelerate ROI for building private cloud on System z with expertise



Product New Features:

- A GTS service designed to assist clients to plan, design, implement and deploy a flexible and robust Linux-based private cloud on IBM z/VM® in IBM System z® environment.

Features / Business Value:

- Provide assessment, planning, education and implementation of Linux-based private cloud on IBM z/VM®.
- Capitalize the benefits of the latest technology of private cloud on secured System z environment to increase flexibility (faster access to computing), scalability, and agility and to reduce costs.
- Deliver processes, tools, skills and best practices that enable a fast and accurate implementation.

Services Benefits:

- **Improve operational efficiency and reduce costs and complexity** by helping you realize the benefits of a private infrastructure as a service (IaaS) cloud and address challenges of operational inefficiency, high capital and operational costs and high risk
- **Increase Return on Investment (ROI)** by helping you build private cloud more rapidly with IBM server and cloud specialists
- **Your IT staff can focus on core business initiatives** by leveraging IBM cloud and server specialists and repeatable methodologies to support a more streamlined deployment of a Linux-based cloud with reduced business disruption.



Learn More: <http://www-935.ibm.com/services/us/en/it-services/server-services.html#Optimize>

IBM Implementation Services für System z - zEnterprise setup and migration



Implement and migrate to zEnterprise environment with IBM experts for better ROI

Product New Features:

Standard service includes:

- Planning, configuration and implementation of the IBM zEnterprise hardware
- Identify potential migration scenarios and access readiness of install
- Identify and order the current zEnterprise, OS (z/OS, z/VM and z/VSE) and subsystem (DB2, CICS, WAS) maintenance
- Review and update hardware, network and storage provisioning. Activities include review and update of I/O configuration, VTAM definition, TCP/IP parm settings, HMC and LPAR configuration, OS (z/OS, z/VM, z/VSE and Linux) configurations
- Assist with Open System Adaptor (OSA) implementation
- Assist with Flash implementation (if applicable)
- One day onsite (customer choice of date within 120 days of contract) to facilitate with deployment

Three optional services that can be acquired individually using on-site daily service assist feature code

- z/OS, subsystem and ISVs migration
- ICB and ISC-3 migration to InfiniBand
- Defined scope of infrastructure perform work specific to client's requirements

Features / Business Value:

- Provide assessment, planning, education and implementation as well as migration to zEnterprise.
- Deliver processes, tools, skills and best practices that enable a fast and accurate implementation and migration.
- Accelerate the adoption of IBM zEnterprise technology and realize the value sooner
- Reasonably priced to ensure success of first deployment

Product Benefits:

- Increase ROI with smooth and faster deployment of IBM zEnterprise leveraging repeatable methodology, automation tools, and project management.
- Minimum business disruption with deployment by IBM server specialists.



Learn More: <http://www-935.ibm.com/services/us/en/it-services/ibm-implementation-services-for-system-z-zenterprise-setup-and-migration.html>

© 2012 IBM Corporation

IBM kann Kunden bei Investitionen in zEnterprise unterstützen



- **Niedrige Nutzungskosten: Fair Market Value (FMV) leases**
- Kosten nur für den Nutzen von IT
 - hilft den Wertverlust bei IT zu verringern
- Kapazitätserhöhungen mit minimalen Budget Veränderungen
- IT erneuern und von TCO profitieren



- **Reduktion finanzieller Risiken mit monatlichen Zahlungen**
- Strukturen um Kosten an Einnahmen anzupassen
- Vorhersehbare Zahlungen versus vorab Kosten
- Verbessert den cash flow mit monatlichen Zahlungen



- **Alternative Finanzierungsgrundlagen um das Budget zu erhalten**
- Fixed rate capital at competitive rates with flexible terms
- Preserve capital for revenue generating initiatives

zEC12



- **5 Modelle**
 - H20, H43, H66, H89 und HA1
- **Prozessor**
 - “hex” (6) Core Prozessor
 - 5.5 Ghz
 - Erweitertes “out of Order” Design
 - 1 Crypto-Coprozessor für jeden Prozessor
 -
- **Prozessor Units (PUs)**
 - 27 (30 für HA1) PU Cores pro Buch
 - Bis zu 16 SAPs pro System, standard
 - 2 Ersatzprozessoren pro System
 - Sub-capacity verfügbar für bis zu 20 CPs
 - 3 sub-capacity Stufen
- **Memory**
 - RAIM Memory design
 - Bis zu 3 TB pro System
 - Flash Express
 - Transactional Memory
 - 2 GB Pages
- **I/O**
 - Bis zu 48 PCIe interconnects pro System à 8 GBps
 - 384 Gigabytes pro Sekunde

Investmentschutz

- **Designed to protect your investment**
 - Offering upgrades from z10 EC™ and z196 to the zEC12
 - Upgrades from zBX Model 002 to zBX Model 003
- **Full upgradeability within the zEC12 family**
 - Upgrade to Model HA1 will require a planned outage
- **On demand offerings offer temporary or permanent growth when you need it**



IBM







Roland Trauner

IBM Sales &
Distribution, STG
Sales

Mainframe Platform, IBM
System z Brand Manager

Vielen Dank



IBM zEnterprise EC12

Operational Analytics – Trusted Security and Resiliency – Efficiency at Scale

Betriebsanalysen

- Bewährte Sicherheit und Ausfallsicherheit

- Effizienz die skaliert

Backup

Flexible Connectivity on zEC12

Improving bandwidth, granularity and options for connections

For Clustering

- **HCA-3 InfiniBand® Coupling Links**
 - 12x InfiniBand
 - 1x InfiniBand (4 ports) **improved scalability**
- **ISC-3 (peer mode only) ¹**
- **STP**
 - **Improved broadband security**

Heterogeneous environments

- **zBX**
 - Intraensemble data network (IEDN)
 - Intranode management network (INMN)
- **To PureSystems**
 - 10 Gb Ethernet

HMC

- **New IBM zAware partition managed from HMC**
- **Location to run Unified Resource Manager – including monitoring CPU, energy, workload performance**
- **Host of the ensemble – controlling all functions of the ensemble**
- **Primary with Alternate needed for DR**

Within zEC12

- **PCIe I/O Infrastructure**
- **I/O Drawer and I/O Cage**
- **Flash Express**



To the Data

- **FICON® Express8S (PCIe-based)**
 - SX and 10KM LX
- **FICON Express4¹**
 - SX and 10 km LX
- **Enhanced channel subsystem**
- **FICON Express8¹**
 - SX and 10KM LX

To the Network

- **OSA-Express4S PCIe based)**
 - 10 Gb Ethernet LR and SR
 - 1 Gb Ethernet SX and LX
- **OSA-Express3¹**
 - 10 Gb Ethernet
 - 1 Gb Ethernet
 - 1000BASE-T Ethernet

¹ Carry forward only

IBM System Storage DS8800

Unique performance, availability and scalability makes DS8800 the ideal storage platform for zEnterprise

- **Self-optimizing performance and cost for hybrid computing**
 - Easy Tier can improve performance by up to 3x by moving only 3% of data to SSDs
 - Integrated QoS management aligns server and storage resources with application priorities
 - 8x faster query performance for operational analytics with special DS8000® List Prefetch Optimizer for High Performance FICON
- **Designed for near-continuous operations with over six-nines availability**
 - Fully-redundant design for near-continuous data access
 - Tight integration between DS8000 remote mirroring and GDPS HyperSwap is designed for over six-nines (99.9999%) availability
- **Optimized for zEnterprise efficiency and scalability**
 - Support for System z Discovery and Auto-Configuration simplifies configuration
 - Extended Address Volumes of up to 1 TB volumes simplifies management of large volumes
 - New support for System z load balancing algorithms can optimize throughput and response times between server and DS8000



Robust tape storage to protect your information

Helping our clients retain data securely and in compliance with regulatory requirements

Comprehensive tape storage product line for System z attachment



**TS1140
Tape Drive**

- Offers high performance (250 MBps) and high native capacity (4 TB) for storage consolidation
- Provides information security with support for encryption and key management
- Supports Write Once Read Many (WORM) cartridges to help satisfy compliance requirements



**TS3500
Tape Library**

- Scalable, automated data retention with up to 2.7 EB capacity with 3:1 compression
- Offers enhanced data availability and reliability with optional dual library accessory
- Provides data security and regulatory compliance via support for tape drive encryption and WORM cartridges



**TS7700
Virtualization
Engine**

- Virtualization solution implements a fully integrated tiered storage hierarchy of disk and tape
- Reduces batch processing time, total cost of ownership and management overhead
- Disk-only models with up to 1.3 PB native tape volume cache
- Grid configurations for information availability and business continuity

Tape is often cost effective versus disk

- Lower price per MB
- Lowest power and cooling storage option available today

Synergy with zEC12 operating systems

z/OS

- Java exploitation of **Transactional Execution** for **increased parallelism and scalability**
- Enhanced security support for **digital signatures**
- Faster problem determination with **IBM zAware for improved availability**
- Improve availability and performance with **Flash Express**
- **2 GB page** support
- Simpler Specialty Engine (**zIIP**) exploitation
- z/OS v1.13 exploitation of new hardware
- Plus **over 4,100** applications enabled on z/OS®

z/VM

- z/VM® **Compatibility support**
- Guest exploitation **support for new encryption technology**
- Improved I/O performance using **High Performance FICON (zHPF)** for guest exploitation

Linux on System z

- **Improved consolidation ratio** through new capacity performance
- Improved I/O performance using **High Performance FICON (zHPF)**
- **Application and Linux optimization** enabled by full exploitation of zArchitecture extensions
- Optimized system setup via **Linux health checker**
- **FCP end-to-end data integrity checking** for applications and storage subsystems
- Plus **over 3,000** applications on System z



z/TPF

- Support for **86 CPUs**
- **Hardware exploitation** for performance improvements

z/VSE

- **64-bit addressing** with z/VSE® V5.1
- Strong **interoperability with Linux** on System z
- New CICS functionality (**CICS Explorer**)

- **AND** with blades on the zBX there are even more options with applications on **AIX, Linux on System x or Microsoft Windows**

Operating System Support for zEC12



- **The following are the minimum operating systems planned to run on zEC12:**

- z/OS
 - z/OS V1.12, V1.13
 - z/OS V1.11, V1.10 Lifecycle Extension
- Linux on System z distributions:
 - SUSE Enterprise Server (SLES) SLES 10 and SLES 11
 - Red Hat Enterprise Linux (RHEL) 6 and RHEL 5
- z/VM
 - VM V5.4, 6.1, 6.2 with PTFs
 - z/VM V6.1, 6.2 for zBX support
- z/VSE
 - z/VSE V4.3, V5.1, with PTFs
 - z/VSE V5.1 with PTFs for Crypto Express4S toleration
- z/TPF V1.1

- **Using the general purpose application server blades we have:**

- Linux: Red Hat RHEL 5.5 and up, 6.0 and up and SUSE Linux Enterprise Server (SLES) 10 (SP4) and up and SLES 11 SP1 and up
- Microsoft Windows Server 2008 R2 and Microsoft Windows Server 2008 (SP2) (for either Windows we recommend Datacenter Edition)
- AIX: AIX 5.3 Technology Level 12 or higher, AIX 6.1 Technology Level 5 or higher, AIX 7.1

IBM zEnterprise EC12 in an information-centric world

Operational analytics to enable business opportunities

- Most sophisticated data warehousing and analytics solutions with the fastest query performance in the market
- Leverage your data to build competitive advantage

Trusted resilience for unmatched security and reliability

- Most secure system with 99.999% reliability
- Unified platform for rapid development of secure applications end-to-end

Efficiency at scale

- Most efficient and fastest system for mission-critical workloads
- Increased efficiency will free up IT resources to focus on new services to drive growth
- Hybrid architecture enables rapid cloud deployment



IBM Global Financing can help clients invest in IBM zEnterprise today



Lower cost of use with Fair Market Value (FMV) leases

- Pay only for use of IT; help reduce IT obsolescence
- Add capacity for minimal budget impact
- Improve asset management with flexible end of lease options
- Refresh IT regularly and improve TCO



Improve cash flow with monthly payments

- Customized structures to better align costs to expected returns
- Predictable payments versus up-front costs



Alternate funding source to help preserve budget

- Fixed rate capital at competitive rates with flexible terms
- Preserve capital for revenue generating initiatives

Redbooks



- **IBM zEnterprise EC12 Technical Introduction, SG24-8050**

This book provides concepts, positioning, and a business value view of zEnterprise System capabilities, hardware functions/features, and associated software support. It is intended for IT Managers, consultants, IT Architects and Specialists, and anyone who wants to understand the basic elements of the zEnterprise EC12.

- **IBM System z Connectivity Handbook, SG24-5444**

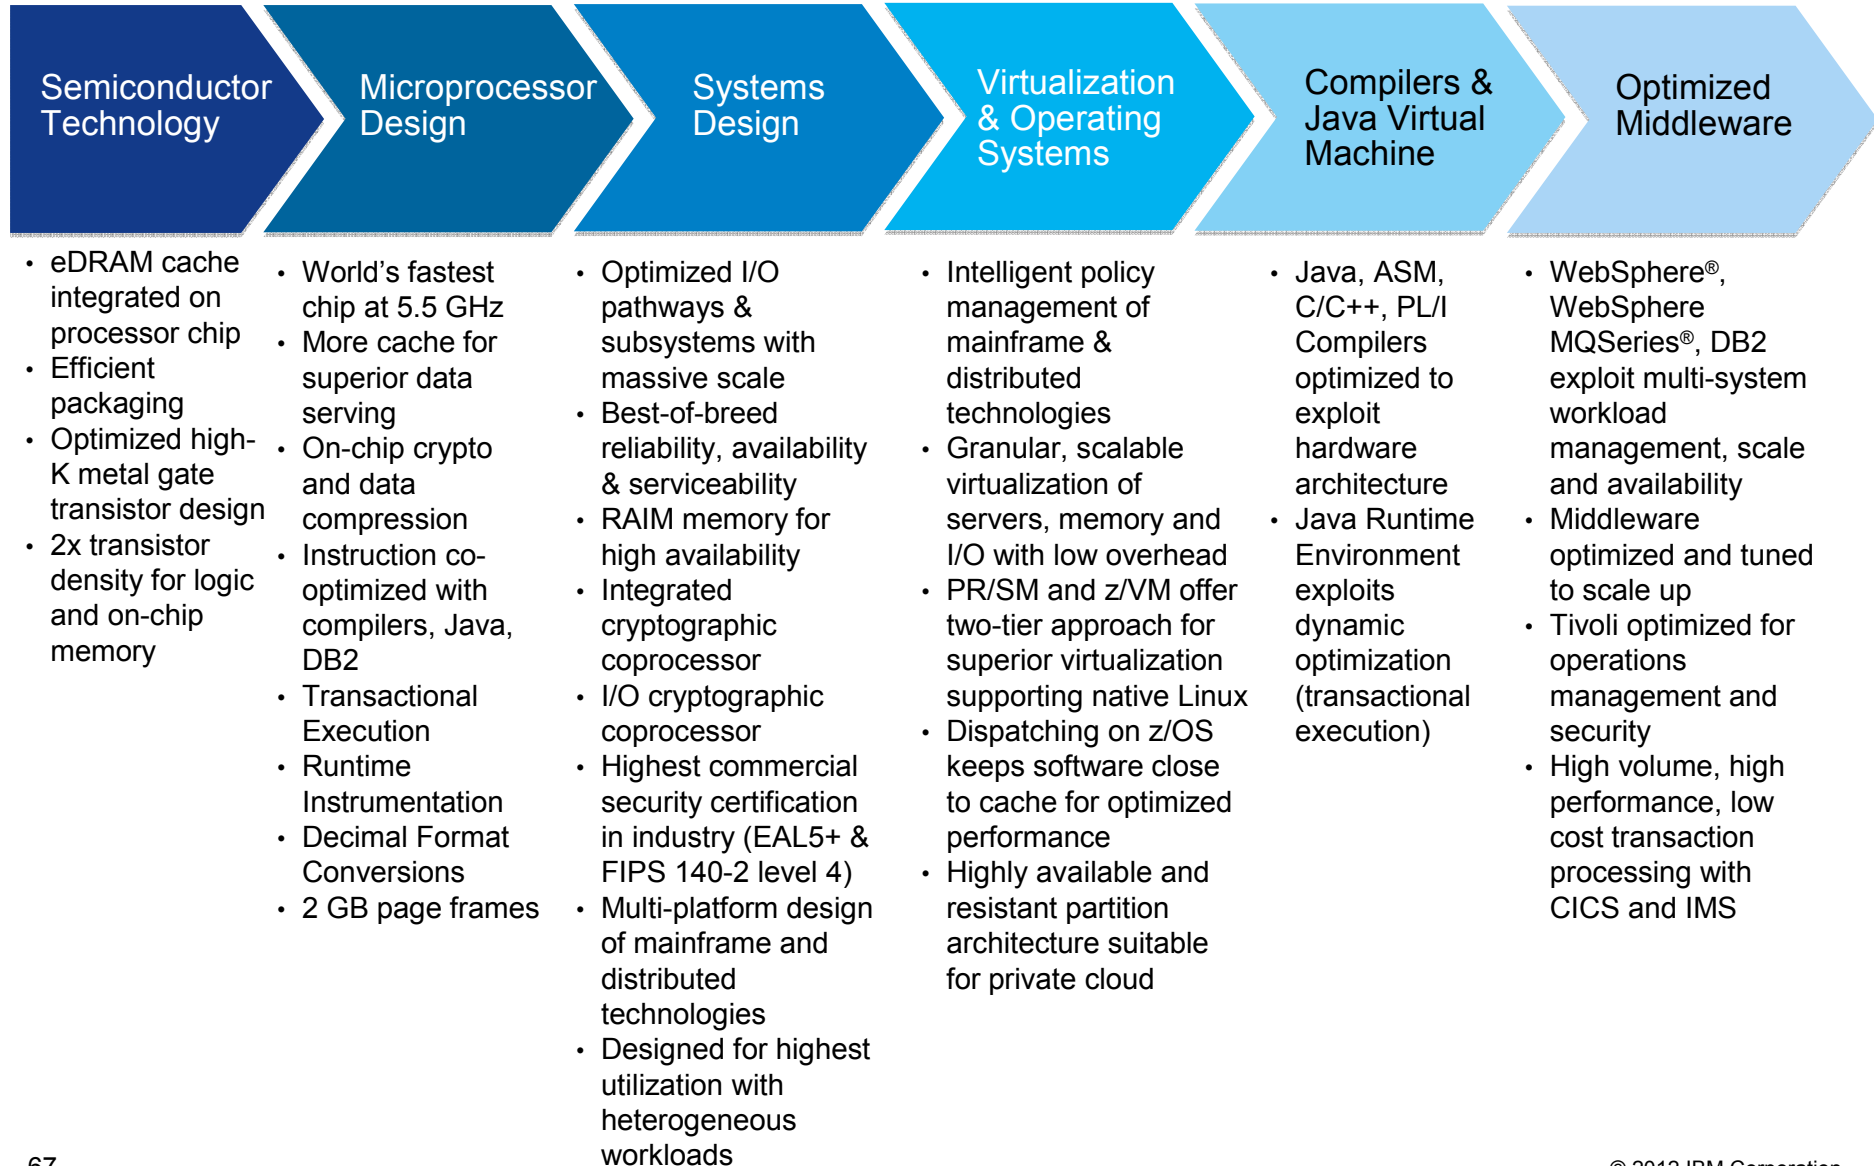
This book highlights the hardware and software components, typical uses, coexistence, and relative merits of the System z I/O features. It is intended for data center planners, IT Specialists, system engineers, technical sales staff, and network planners who are involved in planning connectivity solutions for System z servers.

- **IBM zEnterprise EC12 Technical Guide, SG24-8049**

This book provides specific information about the zEnterprise EC12 (zEC12) and its functions, features, and associated software support. Greater detail is offered in areas relevant to technical planning. It is intended for systems engineers, system programmers (IT Specialists), planners, and anyone wanting to understand the zEC12 functions and plan for their usage.

- <http://www.redbooks.ibm.com/Redbooks.nsf/pages/zEnterprise?Open>

IBM zEnterprise EC12: An optimized system



zEC12 Functional Comparison to z196

Processor / Memory

- Uniprocessor Performance
 - System Capacity
 - Processor Design
 - Cache
 - Models
 - Processing cores
 - Granular Capacity
 - Memory
 - Fixed HSA
- Up to 25% performance improvement over z196 uniprocessor ¹
 - Up to 50% system capacity performance improvement over z196 80-way ¹
 - New 5.4-6¹ GHz processor chip versus 5.2 GHz
 - zEC12 has 33% more L2 cache, instruction and data (total 2 MB versus total 1.5 MB on z196), 100% more L3 cache (total 48 MB versus 24 MB on z196), 100% more L4 cache (384 MB versus 196 on z196)
 - Five models with up to 4 books (z196 had five models)
 - Up to 101 cores to configure, up to 80 on z196
 - Up to 161 capacity settings versus 125 on the z196
 - Up to 3 TB RAIM memory (same as z196)
 - Up to 32 GB fixed HSA versus z196 has 16 GB fixed HSA

Virtualization and Alternative Processors

- Virtualization
 - zEnterprise BladeCenter Extension (zBX)
- zEnterprise Unified Resource Manager provides virtualization management for blades installed in the zBX Mod 003.
 - zEnterprise Unified Resource Manager has “resource workload awareness” where hybrid resources can be managed and optimized across the zEnterprise.
 - zEnterprise System is a truly integrated hardware platform that is able to span and intelligently manage resources across mainframe and distributed technologies – including select POWER7 and IBM System x blades
 - Supported optimizer is DataPower XI50z in the zBX Mod 003.
 - zBX Model 003 (versus zBX Model 002 which attaches to z196)

Connectivity

- HiperSockets™
 - FICON
 - I/O subsystem
 - Internal I/O Bandwidth
 - Coupling
 - Cryptography
- Both zEC12 and z196 support of 32 HiperSockets
 - New OSA-Express4S 1000 BASE-T included in PCIe I/O infrastructure. FICON Express8S and OSA-Express4S adapters available on zEC12, z196, z114
 - zEC12 has industry standard 8 Gbps InfiniBand supports high speed connectivity and high bandwidth
 - Coupling with HCA-3 InfiniBand Coupling Links
 - Crypto Express4S enhanced with new FIPS 140-2 Level 4 cert and PKCS#11 support
 - Elliptic Curve Cryptography (ECC)

RAS

- RAS Focus
 - Availability
- New IBM zAware offers high speed analytics facilitates the ability to consume large quantities of message logs for smarter monitoring
 - zEC12 offers advanced memory enhancements (RAIM) and advanced power and thermal optimization and management that can help to control heat / improve RAS
 - New PCIe Flash Express on zEC12 to handle paging workload spikes and improve availability – not available on z196

Environmentals

- Energy
 - Cooling
- Power Save modes for processor
 - New improved integrated cooling system
 - Optional Non Raised Floor and overhead cabling options for both I/O and (New!) Power
 - Optional water cooling and DC power

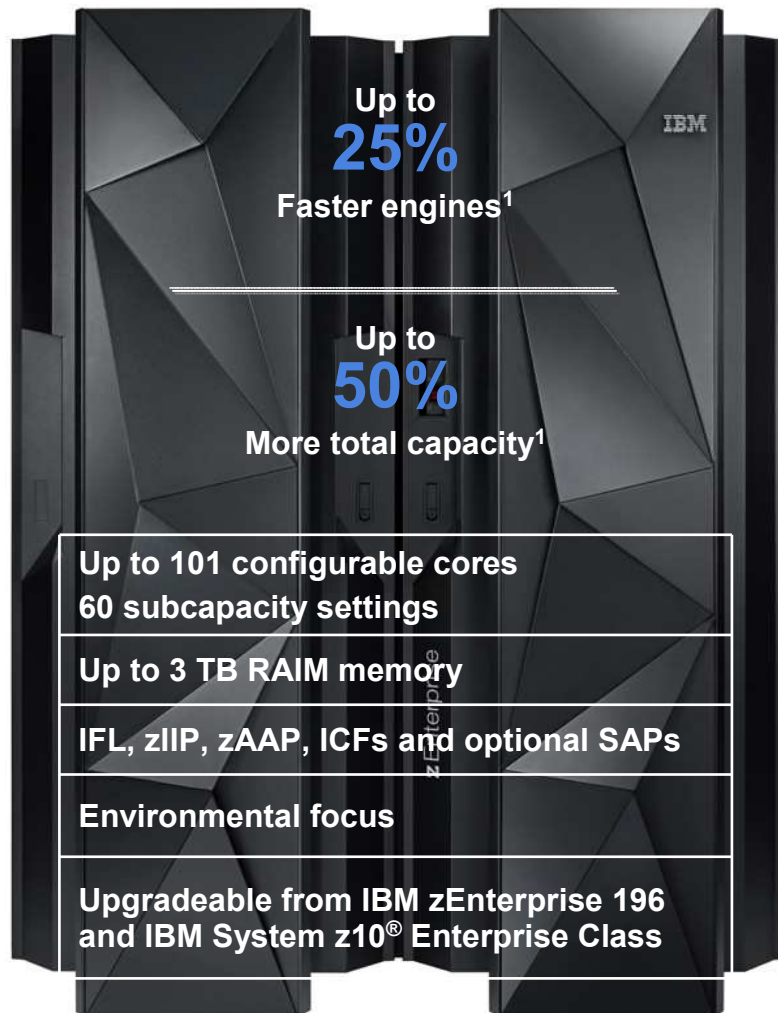
¹ Based on preliminary internal measurements and projections. Official performance data will be available upon announce and can be obtained online at LSPR (Large Systems Performance Reference) website at: <https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lspindex?OpenDocument>. Actual performance results may vary by customer based on individual workload, configuration and software levels.

zEC12 Overview

- **Machine Type**
 - 2827
- **5 Models**
 - H20, H43, H66, H89 and HA1
- **Processor Units (PUs)**
 - 27 (30 for HA1) PU cores per book
 - Up to 16 SAPs per system, standard
 - 2 spares designated per system
 - Dependant on the hardware model - up to 20, 43, 66, 89, 101 PU cores available for characterization
 - Central Processors (CPs), Internal Coupling Facility (ICFs), Integrated Facility for Linux (IFLs), System z Application Assist Processors (zAAPs), System z Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs)
 - Sub-capacity available for up to 20 CPs
 - 3 sub-capacity points
- **Memory**
 - RAIM Memory design
 - System Minimum of 32 GB
 - Up to 768 GB per book
 - Up to 3 TB for System and up to 1 TB per LPAR
 - 32 GB Fixed HSA, standard
 - 32/64/96/112/128/240/256 GB increments
 - Flash Express
- **I/O**
 - 6 GBps I/O Interconnects – carry forward only
 - Up to 48 PCIe interconnects per System @ 8 GBps each
 - Up to 4 Logical Channel Subsystems (LCSSs)
 - Up to 3 Sub-channel sets per LCSS
- **STP - optional (No ETR)**



zEnterprise EC12 is the core of next generation System z



zEC12

Machine Type: 2827

Models: H20, H43, H66, H89, HA1

Advanced Technology 5.5 GHz processor chip for performance boost for all workloads

- Over **78,000 MIPS** for large scale consolidation
- **Larger cache** for data serving

Processor chip optimized for software performance – exploited by *Java, PL/I, compilers, DB2* and more

Innovation to drive availability to superior levels

- **IBM zAware** offers snap-shot of the current state of your business
- **FLASH Express and pageable large pages** to drive availability and performance for critical workloads


Trusted resilience is a zEnterprise standard

- High speed **cryptography integrated as part of the chip**
- Enhanced functions with new **Crypto Express4S**
- PR/SM designed for **EAL5+ certification**

¹Based on preliminary internal measurements and projections against a z196. Official performance data will be available upon announce and can be obtained online at LSPR (Large Systems Performance Reference) website at: <https://www.ibm.com/servers/resourceLink/lib03060.nsf/pages/lspindex>. Actual performance results may vary by customer based on individual workload, configuration and software levels.

zEC12 – Optimization and scale improvements starting at the core

- **New 5.5 GHz 6-Core Processor Chip continues our leadership in microprocessor design with a boost in performance for all workloads**
 - Second generation out of order execution design
- **Larger caches to optimize data serving environments**
- **New hardware functions optimized for software performance**
 - ***Transactional Execution Facility*** for parallelism and scalability
 - ***Runtime Instrumentation Facility*** is intended to help reduce Java overhead
 - ***2 GB page frames*** are intended to offer performance improvements for DB2 buffer pools and Java heaps
 - New IBM Enterprise PL/I compiler is planned to exploit and get a performance boost from ***decimal format conversions facility***
- **Integrated cryptographic function available on each core characterized as a CP or IFL with enabling microcode**



zEC12 Processor Unit allocation/usage

Model	Books/ PUs	CPs	IFLs uIFLs	zAAPs	zIIPs	ICFs	Std SAPs	Optional SAPs	Std. Spares	Rsvd. PUs
H20	1/27	0-20	0-20 0-19	0-10	0-10	0-20	4	0-4	2	1
H43	2/54	0-43	0-43 0-42	0-21	0-21	0-43	8	0-8	2	1
H66	3/81	0-66	0-66 0-65	0-33	0-33	0-66	12	0-12	2	1
H89	4/108	0-89	0-89 0-88	0-44	0-44	0-89	16	0-16	2	1
HA1	4/120	0-101	0-101 0-100	0-50	0-50	0-101	16	0-16	2	1

- zEC12 Models H20 to H89 use books with 27 core MCMs. The Model HA1 has 4 books with 30 core MCMs
 - Each MCM uses PU chips with a combination of 4, 5 and 6 active cores
- The maximum number of logical ICFs or logical CPs supported in a CF LPAR is 16
- The Reserved PU is not available for customer purchase
- Concurrent Book Add is available to upgrade from model H20 to model H89

Notes: 1. At least one CP, IFL, or ICF must be purchased in every machine
 2. One zAAP **and** one zIIP may be purchased for each CP purchased even if CP capacity is “banked”.
 3. “uIFL” stands for Unassigned IFL

zBX Overview



- **Machine Type/Model 2458-003**
- **Racks – Up to 4 (B, C, D and E)**
 - 42U Enterprise, (36u height reduction option)
 - 4 maximum, 2 chassis/rack
 - 2-4 power line cords/rack
 - Non-acoustic doors as standard
 - Optional Acoustic Doors
 - Optional Rear Door Heat Exchanger (conditioned water required)
- **Chassis – Up to 2 per rack**
 - 9U BladeCenter
 - Redundant Power, cooling and management modules
 - Network Modules
 - I/O Modules
- **Blades (Maximum 112 single width blades in 4 racks)**
 - Customer supplied POWER7 Blades (0 to 112)
 - Customer supplied IBM System x Blades (0 to 56)
 - DataPower XI50z, M/T 2462-4BX (0 to 28 – double width)
- **Management Firmware**
 - Unified Resource Manager
- **Top of Rack (TOR) Switches - 4**
 - 1000BASE-T intranode management network (INMN)
 - 10 GbE intraensemble data network (IEDN)
 - GbE IEDN for customer network
- **Network and I/O Modules in the BladeCenter**
 - 1000BASE-T and 10 GbE modules
 - 8 Gb Fibre Channel (FC) connected to customer supplied disks

zBX offers Flexibility for Workload Deployment and Integration

- **zBX is built with integrated IBM certified components**
 - Standard parts – TOR switch, BladeCenter Chassis, Power Distribution Units, Optional Acoustic Panels
 - Optional optimizer - IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z) ordered as a feature of zBX
- **Up to 112 blades are supported on zBX**
 - System x and POWER7 blades are acquired through existing channels
 - IBM System x Blades – up to 56
 - IBM BladeCenter HX5 (7873) dual-socket 16-core blades, four supported memory configurations for zBX – 64 GB, 128 GB, 192 GB, 256 GB
 - IBM POWER7 Blades – up to 112
 - IBM BladeCenter PS701 Express - 8-core processor 3.0GHz, three configurations supported in zBX - 32 GB, 64 GB, 128 GB
 - AIX: AIX 5.3 Technology Level 12 or higher, AIX 6.1 Technology Level 5 or higher, AIX 7.1
 - Up to 28 DataPower XI50z blades (double wide)
 - Ability to mix and match blades in the same chassis and number of blades supported varies by type
- **System z support – Blades assume System z warranty and maintenance when installed in the zBX**
- **Investment protection**
 - Upgrade the Model 002 to Model 003

IBM zEnterprise
BladeCenter Extension (zBX)
Machine Type: 2458 Mod 003
(for attachment to zEC12)



Operating System Environments extend application flexibility

- **Operating Systems are customer acquired and installed**
- **Unified Resource Manager will install hypervisor on blades in the zBX**
 - Integrated hypervisor (KVM-based) for System x blades
 - PowerVM™ Enterprise Edition for POWER7 blades
- **Support for Linux and Windows environments on System x blades in zBX**
 - 64-bit version support only
 - Linux: Red Hat RHEL 5.5 and up, 6.0 and up & SUSE Linux Enterprise Server (SLES) 10 (SP4) and up and SLES 11 SP1 and up
 - Microsoft Windows Server 2008 R2 and Microsoft Windows Server 2008 (SP2) (for either Windows we recommend Datacenter Edition)
- **Support of AIX environments on POWER7 blades in zBX**
 - AIX: AIX 5.3 Technology Level 12 or higher, AIX 6.1 Technology Level 5 or higher, AIX 7.1
- **Certifications inherited from blades**
 - SAP support for Linux and Windows on x86 blades in the zBX
- **PowerHA™ SystemMirror Standard Edition for AIX supported for the zBX with PS701 blades**



System z software addressing key business initiatives

Transactional Processing

Millions of transactions per day

Business Rules and Processes

Create reusable services and processes

Security

Advanced encryption and decryption software

Enterprise Modernization

Single platform for development

Integrated service management

Single point of control

Integrated Appliances

Packaged solutions for optimal performance

Cloud Computing

SaaS solutions

Data Warehousing


Unique temporal capabilities

Operational Analytics

Secured right-time analysis

Social Business

Integrating social business with enterprise computing



... an undeniable foundation for solutions

SUSE Offering Summary *

for SUSE Linux Enterprise Server for System z

Attractive subscriptions and services are available: talk to us!

- **Multi year, multiple IFL subscriptions offer monetary value – scale & save**
 - For larger volumes, you get up to 45% of discount per subscription
 - <http://www.suse.com/products/systemz/how-to-buy>
- **SUSE Linux Enterprise High Availability Extension included – prevent outages**
 - Additionally, save more than \$1.800 per IFL compared to distributed systems
 - http://www.suse.com/products/systemz/features/linux_server_clustering.html
- **zBX offering – 1 IFL, multiple zBX subscriptions – boost consolidation**
 - Save up to \$39k running your ensemble with SUSE Linux Enterprise Server
 - <http://www.suse.com/promo/zbx.html>
- **Long Term Support Service available – preserve environments**
 - <http://www.suse.com/support/programs/long-term-service-pack-support.html>



SUSE Multi IFL Pricing USD

Price Listed Per IFL

Basic Subscription Pricing*

	1 IFL	2-5 IFLs	6-11 IFLs	12-25 IFLs	26-53 IFLs	54+ IFLs
1 year	\$11,999	\$10,199	\$9,599	\$8,999	\$7,799	\$6,599
3 year	\$32,400	\$27,538	\$25,918	\$24,298	\$21,058	\$17,819
5 year	\$48,000	\$40,797	\$38,397	\$35,997	\$31,197	\$26,398

Priority Subscription Pricing*

	1 IFL	2-5 IFLs	6-11 IFLs	12-25 IFLs	26-53 IFLs	54+ IFLs
1 year	\$18,000	\$15,300	\$14,400	\$13,500	\$11,700	\$9,900
3 year	\$48,600	\$41,310	\$38,880	\$36,450	\$31,590	\$26,730
5 year	\$72,000	\$61,200	\$57,600	\$54,000	\$46,800	\$39,600

Basic is tailored to be combined with IBM GTS services.

Multi-IFL pricing information is located [online here](#).

Get more details on Basic or Priority service levels [here](#). (May 2012, prices are subject to change)0

Red Hat Offering Summary

RED HAT ENTERPRISE LINUX

The world's leading open source application platform on one certified platform

- Applications - Thousands of certified ISV applications
- Deployment - Including standalone or virtual servers, cloud computing, or software appliances
- Hardware - Wide range of platforms from the world's leading hardware vendors

Comprehensive enterprise infrastructure

- Server Virtualization
- Storage Virtualization
- Systems management, provisioning, monitoring

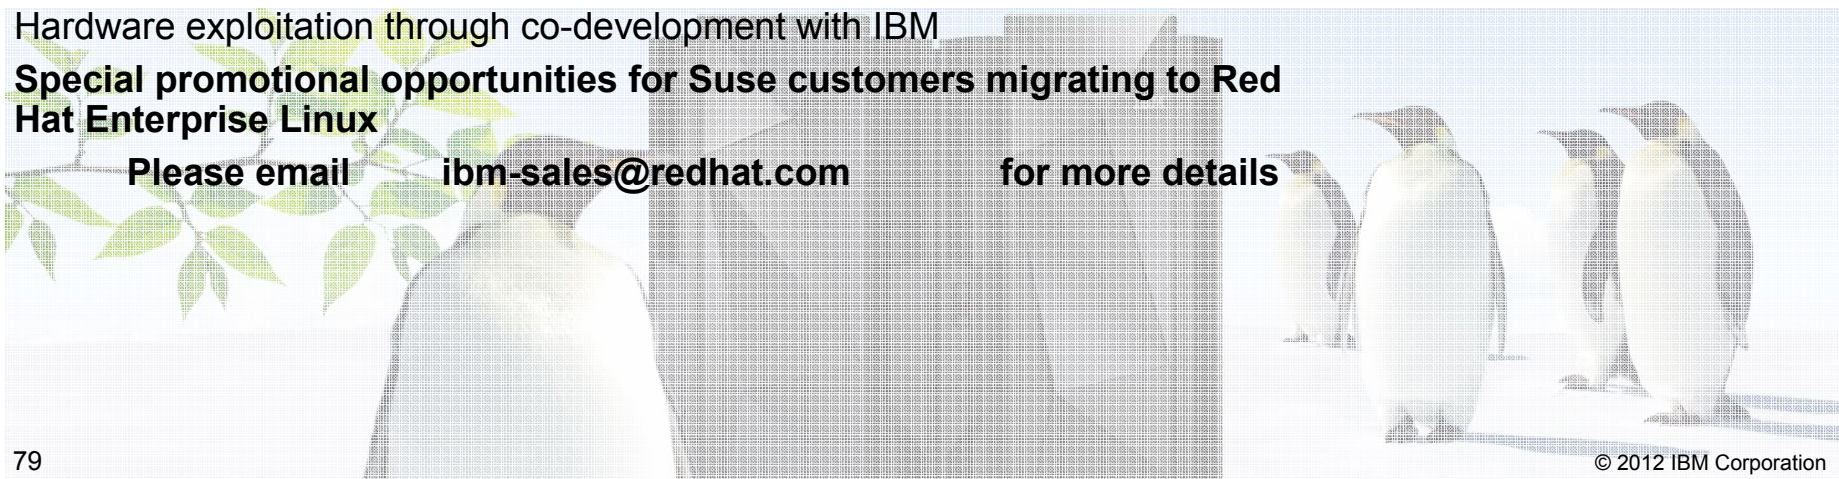
Push TCO by using a single Operating System from Laptop to Mainframe

Highest security accreditations (via SELinux)

Hardware exploitation through co-development with IBM

Special promotional opportunities for Suse customers migrating to Red Hat Enterprise Linux

Please email ibm-sales@redhat.com for more details



IBM POWER7 and System x Blades

General purpose processors under one management umbrella

What is it?

The zBX infrastructure can host select IBM POWER7 and IBM System x blades. Each blade comes with an installed hypervisor that offers the possibility of running an application that spans z/OS, Linux on System z, AIX on POWER®, Linux or Microsoft® Windows® on System x but have it under a single management umbrella.



IBM BladeCenter PS701



IBM BladeCenter HX5 (7873)

How is it different?

- **Complete management:** Advanced management brings operational control and cost benefits, improved security, workload management based on goals and policies.
- **Virtualized and Optimized:** Virtualization means fewer resources are required to meet peak demands with optimized interconnection. Multiple resources (both blade types and optimizers) can reside in a single zBX.
- **Integrated:** Integration with System z brings heterogeneous resources together that can be managed as one.
- **Transparency:** Applications certified to run on AIX 5.3, 6.1 or 7.1 on POWER7 blades and those certified to run on supported releases of Linux on System x or Windows on the System x blades will run on those blades in a zBX. No changes to deployed guest images.
- **More applications:** Brings larger application portfolio to System z.

IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise Helps Extend the Value of zEnterprise

Purpose-built hardware for simplified deployment and hardened security helps businesses quickly react to change and reduce time to market

What is it?

The IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise can help simplify, govern, secure and integrate XML and IT services by providing connectivity, gateway functions, data transformation, protocol bridging, and intelligent load distribution.

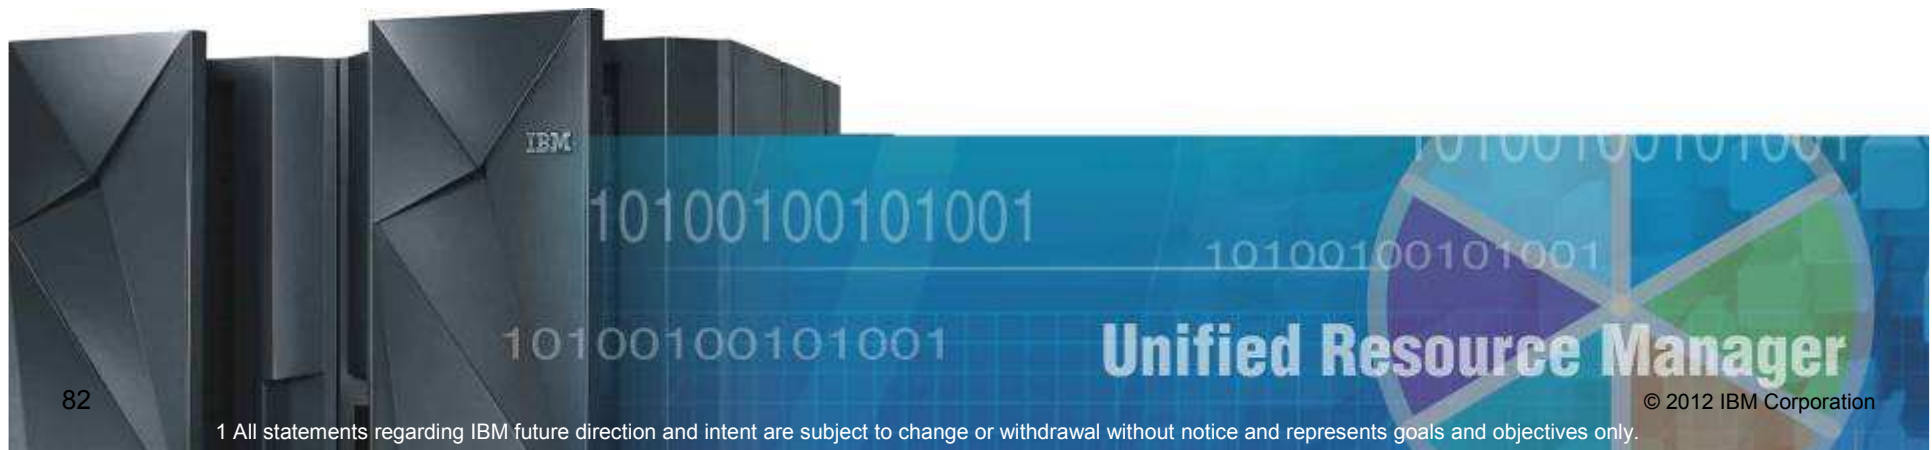


How is it different?

- **Security:** VLAN support provides enforced isolation of network traffic with secure private networks.
- **Improved support:** Monitoring of hardware with “call home” for current/expected problems and support by System z Service Support Representative.
- **System z packaging:** Increased quality with pre-testing of blade and zBX. Upgrade history available to ease growth.
- **Operational controls:** Monitoring rolled into System z environment from single console. Consistent change management with Unified Resource Manager.
- **Cloud:** WebSphere DataPower enhancements can provide a secure, managed connection from the enterprise applications or enterprise users to public cloud applications.

Unified Resource Manager optimizes system resources

- Bringing mainframe governance for System z resources
- Single view of virtualized resources across platforms
- Integrated network for better security, control and faster time to value
- Management of resources as defined by your business goals and objectives
- Programmable interfaces (APIs) to connect with system management tools for total management capabilities
 - IBM Tivoli Monitoring (ITM) has been enhanced to use the APIs
- New dynamic storage capabilities for System x blades



zEnterprise Unified Resource Manager Pricing Strategy

Priced to Value: Tiered functionality that scales

- **Manage, Advanced Management and Automate are ordered as features of the zEC12**


- When placing an order for a zEC12, the default is 'Manage' FC#0019. If you want "Advanced Management / Automate" you will need to order FC#0020. The additional feature codes for these functions for the blades/optimizers will automatically be generated for you based on the number of blades/optimizers that you order. To get ensemble management and cables make sure that you also order FC#0025 on the zEC12.

Manage	Delivers Unified Resource Manager's function for core operational controls, installation and configuration, and energy monitoring.
Advanced Management	Delivers workload definition and performance policy monitoring and reporting.
Automate	Delivers workload definition and performance policy monitoring and reporting. Delivers goal oriented monitoring and goal oriented management of resources and energy management.

	Manage – <u>per connection</u>	Advanced Management - <u>per connection</u>	Automate – <u>per connection</u>
zEC12 base hardware configuration	FC#0019 - N/C	N/A	FC#0020 - N/C
IFL	N/C	N/A	FC#0054 – Yes
POWER7 Blade (zBX FC#0612)	FC#0048 – Yes	N/A	FC#0051 – Yes
DataPower Blade (zBX FC#0611)	FC#0047 – Yes	N/A	FC#0050 – N/C
System x Blade (zBX FC#0613)	FC#0049 – Yes	FC#0053 – Yes	N/A

Parallel Sysplex and Coupling Facility improvements *CFCC Level 18*

- **Scalability enhancements**
 - Up to 64 1x IFB3 links
 - 16 logical engines per CF LPAR
- **RMF™ channel path details**
 - Differentiate various IFB link types
 - Detect if CIB link running "degraded"
- **Serviceability enhancements**
 - Additional structure control info in CF dumps
 - Enhanced triggers for CF non-disruptive dumping.
 - Enhanced CFCC tracing support
- **Performance enhancements**
 - DB2 GBP cache bypass
 - Dynamic structure size alter improvement
 - Cache structure management



Power 7+ Ankündigung vom 3. Oktober

Die neuen Enterprise Power Systems

Neue Power Systems Hardware

- Neue Power 770, 780 und verbesserte 795
- POWER7+ Prozessoren & Architektur
- Elastic Capacity (CoD) für Power System Pools

Neue Power Systems Software

- Verbesserte PowerVM Virtualisierung
- PowerVM Virtualization Performance Advisor
- PowerVM Live Partition Mobility Erweiterungen
- PowerHA HyperSwap

Neue Power Systems Offerings

- Power Systems Solution Edition für Cloud
- AIX Solution Editions für Cognos und SPSS

Neue Power Systems Services

- Neue GTS Power Services



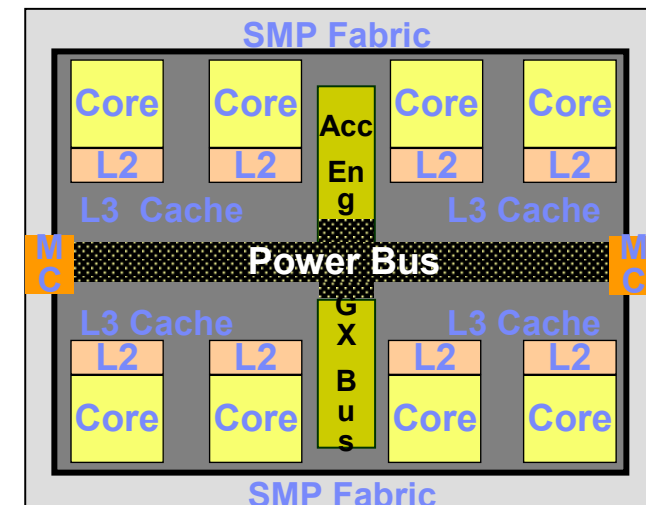

Einführung POWER7+

Physisches Design:

- Integrierter Cache, Memory Controllers und Akzeleratoren
- 32nm Technologie

Features:

- ✓ Höhere Frequenzen
- ✓ Größerer L3 Cache
- ✓ Memory Compression Accelerator
 - Active Memory Expansion
- ✓ Hardware Encryption Support für AIX
- ✓ Random Number Generator
- ✓ Verbesserte RAS features
- ✓ Verbessertes Energy / Power Gating
- ✓ Verbesserte GX System Buses
- ✓ Verbesserte Single Precision Floating Point Performance
- ✓ 20 virtuelle Maschinen pro Kern

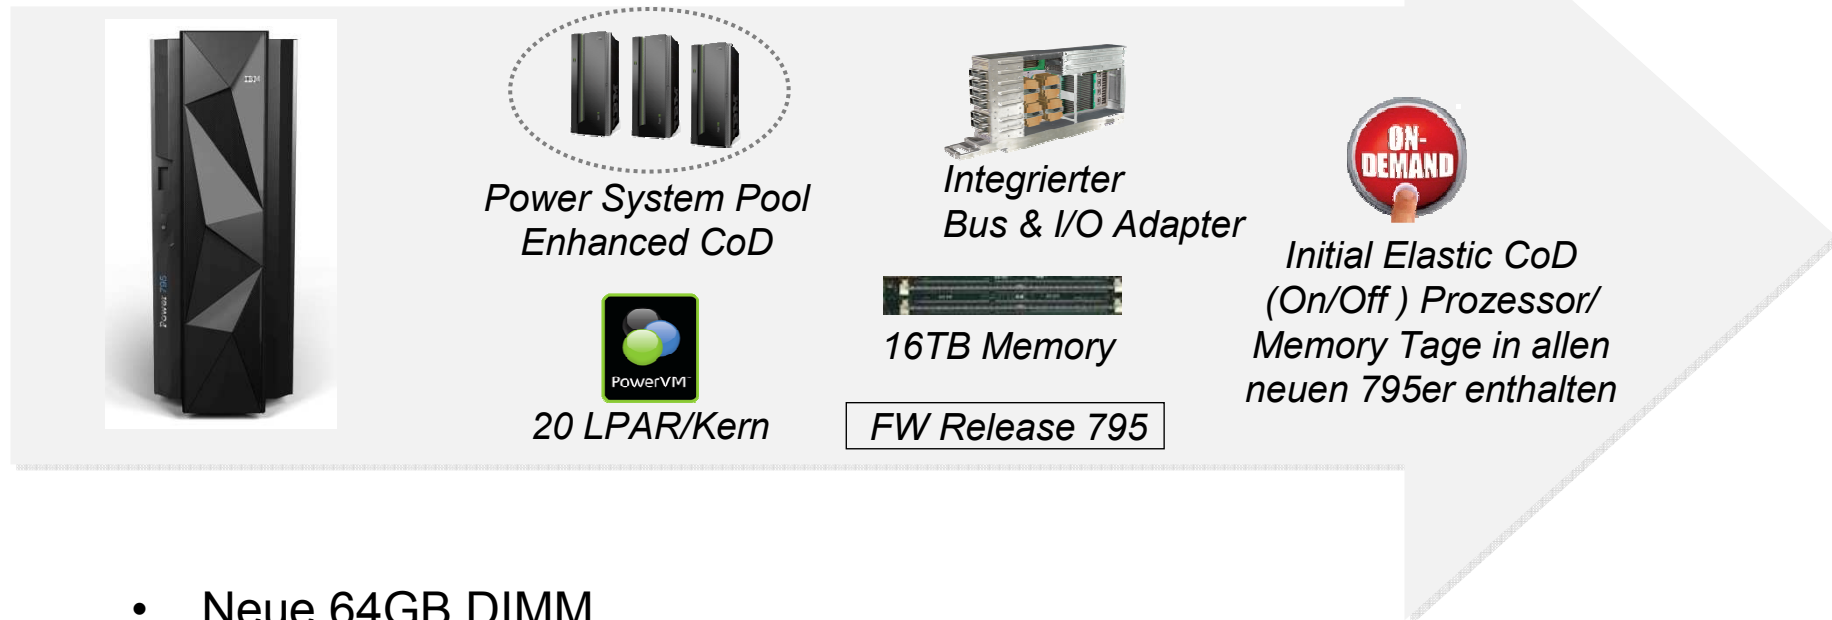


POWER7+
32 nm

Power Systems 770, 780

Benefit	Neue Power 770	Neue Power 780
Mehr Leistung	• POWER7+	
	• 10 MB L3 cache pro Kern	
	- 64 Kerne @ 3.8 GHz - 48 Kerne @ 4.2 GHz	- 128 Kerne @ 3.72 GHz - 64 Kerne @ 4.42 GHz
	• Memory Komprimierung • AIX Dateisystemverschlüsselung	
	• Random Number Generator • Verbesserte GX System Busse • Verbesserte Single Precision Floating Point performance	
	• Dynamic Platform Optimizer	
Höhere Verfügbarkeit	• Self-healing Fähigkeit für L3 Cache Funktionen • Dynamic processor fabric bus repair • Processor Reinitialisierung	
Energieeffizienter	• Verbessertes Energy / Power Gating	
Verbesserte Skalierbarkeit	• 20 VMs pro Kern • 16 gleichzeitige LPM Transfers	
	• Einfachere CoD Aktivierung	• Inkludierte Elastic CoD (On/Off) Prozessor und Memory Tage • Einfachere CoD Aktivierung

Verbesserung der Power 795



- Neue 64GB DIMM
 - ermöglicht der IBM Power 795 bis zu 16TB Hauptspeicher
- Neue Hybrid I/O Adapter
- No-charge Elastic CoD (On/Off) Prozessor
- 20 VMs pro Core

Happy Family

