



The CPU-Measurement Facility Extended Counters Definition for z10, z196/z114, zEC12/zBC12, z13/z13s, z14 and z15

January 22, 2020

SA23-2261-06

Note: Before using this information and the product it supports, be sure to read the general information under “Notices” on page 4.

Sixth Edition (January 2020)

This edition obsoletes and replaces *The CPU-Measurement Facility Extended Counters Definition for z10, z196/z114,zEC12, z13/z13s, z14 and z15 (SA23-2261-05)*.

January 22, 2020

The facilities discussed in this publication are available on certain IBM z Systems™ Processor Complexes. The information published herein should not be construed as implying any intention by IBM to provide these facilities on models other than those described herein.

This publication is provided for use in conjunction with other relevant IBM publications, and IBM makes no warranty, express or implied, relative to its completeness or accuracy. The information in this publication is current as of its publication date but is subject to change without notice.

© Copyright International Business Machines Corporation 2008-2020. All rights reserved.

US Government Users Restricted Rights — Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	4
Trademarks	5
Preface	6
Extended Counters Definitions for z10	7
Extended Counters Definitions for z196 and z114 ..	9
Extended Counters Definitions for zEC12 and zBC12.....	11
Extended Counters for z13 and z13s.....	14
Extended Counters for z14.....	18
Extended Counters for z15.....	22

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

Notices

IBM Corporation
New Orchard Rd.
Armonk, NY 10504
U.S.A.

Produced in the United States of America
January 22, 2020
All Rights Reserved

The information in this document is intended to provide guidance for those implementing CPU-measurement architecture. It discusses findings based on a solution that was created and tested under laboratory conditions. These findings may not be realized in all customer environments, and implementation in such environments may require additional steps, configurations, and performance analysis. The information herein is provided "AS IS" with no warranties, express or implied. This information does not constitute a specification or form part of the warranty for any IBM product. Implementation and certification of the solution rests on the implementation team. The users of this document should always check the latest release information for the applicable product and check the product Web pages for the latest updates and findings.

References in this publication to IBM products or services do not imply that IBM intends to make them available in every country in which IBM operates. Consult your local IBM business contact for information on the products, features and services available in your area.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY, 10504-1785 USA.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models. This equipment is subject to all applicable FCC rules and will comply with them upon delivery. Information concerning non-IBM products was obtained from the suppliers of those products. Questions concerning those products should be directed to suppliers.

Trademarks

IBM, the IBM logo, and z/Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

IEEE is a trademark of the Institute of Electrical and Electronics Engineers, Inc. in the United States, other countries, or both.

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

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

Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems Inc. in the United States and other countries.

Other trademarks and registered trademarks are the properties of their respective companies.

Preface

This document describes what the Extended Counters of the CPU-Measurement Facility count for the z10, z196, z114, zEC12, zBC12, z13, z13s, z14 and z15 families of machines.

Documentation referenced in this publication include the following:

1. *The Set-Program-Parameter and CPU-Measurement Facilities* (SA22-2260), IBM Corp. 2008.
2. *Design and Microarchitecture of the IBM System z10 Microprocessor*, C.-L. K. Shum, et al., “The IBM Journal of Research and Development,” Volume 53, Number 1, November 2009
3. *IBM System z10 Processor Cache Subsystem Microarchitecture*, P. Mak, et al. “The IBM Journal of Research and Development,” Volume 53, Number 1, November 2009.
4. *IBM zEnterprise 196 Microprocessor and Cache Subsystem*, F Busaba, et al., “The IBM Journal of Research and Development,” Volume 56, Number 1.2, January-February 2012
5. *IBM z/Architecture Principles of Operation* (SA22-7832-09), September, 2012.
6. *IBM z13 multithreaded microprocessor*, B. W. Curran, et al., “The IBM Journal of Research and Development” Volume 59, Number 4/5, July/September 2015
7. *IBM z13 processor cache subsystem*, C. R. Walters, et al., “The IBM Journal of Research and Development” Volume 59, Number 4/5, July/September 2015
8. *Performance innovations in the IBM z14 platform*, C. R. Walters, et al., “The IBM Journal of Research and Development” Volume 62, Issue 2/3, March-May 2018”
9. *Design of the IBM z14 microprocessor*, C. Jacobi, et al., “The IBM Journal of Research

Extended Counters Definitions for z10

The Extended Counters provide information about hardware facilities and structures that are specific to a machine family. What a particular Extended Counter number counts can change between machine families. Software can not expect these counters to remain unchanged when moving from one machine family to another. The z10 family includes both EC and BC class machines and these Extended Counter definitions are the same for both.

The Extended Counters are designed to expand upon information provided by the Basic Counter Set that is described in the Set-Program-Parameter and CPU-Measurement Facilities, SA23-2260. The Extended Counter Set provides data that is substantially accurate and is suitable for software performance analysis and tuning. However, under certain cases, the values may not precisely represent the behavior of the CPU and the system activity.

Extended Counters 128 to 135 provide information about events in the cache hierarchy of the z10 family of machines beyond the Level-1 caches that the Basic Counter Set covers. These counters provide information about where a Level-1 cache directory write found its associated cache line. These counters provide a break out of where the data that was returned to Level-1 cache was sourced from.

The Extended Counters 136 to 142 and 145 to 147 contain additional counters related to the cache hierarchy and processor.

For z10 two sets of identifiers have been associated with the various cache levels. The formal definition assigns integer values to each level in the hierarchy, while a colloquial definition allows fractional values (e.g. L1.5). For completeness, both identifiers are provided here. The formal (integer) value is provided along with the informal value in parenthesis, e.g. Level-2 (L1.5).

Counter 128 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from the Level-2 (L1.5) cache.

Counter 129 – A directory write to the Level-1 Data cache directory where the installed cache line was sourced from the Level-2 (L1.5) cache.

Counter 130 – A directory write to the Level-1 Instruction cache directory where the installed cache lined was sourced from the Level-3 cache that is on the same book as the Instruction cache (Local L2 cache).

Counter 131 – A directory write to the Level-1 Data cache directory where the installed cache line was sourced from the Level-3 cache that is on the same book as the Data cache (Local L2 cache).

Counter 132 – A directory write to the Level-1 Instruction cache directory where the installed cache line was sourced from a Level-3 cache that is not on the same book as the Instruction cache (Remote L2 cache).

Counter 133 – A directory write to the Level-1 Data cache directory where the installed cache line was sourced from a Level-3 cache that is not on the same book as the Data cache (Remote L2 cache).

Counter 134 – A directory write to the Level-1 Data cache where the installed cache line was sourced from memory that is attached to the same book as the Data cache (Local Memory).

Counter 135 – A directory write to the Level-1 Instruction cache where the installed cache line was sourced from memory that is attached to the same book as the Instruction cache (Local Memory)

Directory writes that are sourced from memory that is attached to a different book than the book where the Level-1 cache is (Remote Memory) can be back-calculated by using Counter 2 Level-1 I-Cache Directory Write Count and Counter 4 Level-1 D-Cache Directory Write Count that are part of the Basic Counter Set in the CPU-Measurement facility.

Directory writes to the Level-1 Instruction cache that was sourced from Remote Memory =

Counter 2 – (Counter 128 + 130 + 132 + 135)

Directory writes to the Level-1 Data cache that was sourced from Remote Memory =

Counter 4 – (Counter 129 + 131 + 133 + 134)

Counter 136 – A directory write to the Level-1 Data cache where the line was originally in a Read-Only state in the cache but has been updated to be in the Exclusive state that allows stores to the cache line.

Counter 137 – A cache line in the Level-1 Instruction cache has been invalidated by a store on the same CPU as the Level-1 Instruction cache.

Counter 138 – A translation entry has been written into the Level-1 Instruction Translation Lookaside Buffer (ITLB1).

Counter 139 – A translation entry has been written to the Level-1 Data Translation Lookaside Buffer (DTLB1).

Counter 140 – A translation entry has been written to the Level-2 TLB Page Table Entry arrays.

Counter 141 – A translation entry has been written to the Level-2 TLB Combined Region Segment Table Entry arrays.

Counter 142 – A translation entry has been written to the Level-2 TLB Combined Region Segment Table Entry arrays for a one-megabyte large page translation.

Counter 143 - Undefined.

Counter 144 – Undefined.

Counter 145 – Level-1 Instruction TLB miss in progress. Incremented by one for every cycle an ITLB1 miss is in progress.

Counter 146 – Level-1 Data TLB miss in progress. Incremented by one for every cycle a DTLB1 miss is in progress.

Counter 147 – Incremented by one for every store sent to Level-2 (L1.5) cache.

Counter 148 – Undefined.

Counter 149 – Undefined.

Counter 150 – Undefined.

Counter 151 – Undefined.

Extended Counters Definitions for z196 and z114

The Extended Counters provide information about hardware facilities and structures that are specific to a machine family. While much of the information provided in the Extended Counters for z196 and z114 are similar in kind to z10, details of the machines' micro-architectures differ and the z196/z114 Extended Counter definitions are not the same as z10.

The Extended Counters are designed to expand upon information provided by the Basic Counter Set that is described in the Set-Program-Parameter and CPU-Measurement Facilities, SA23-2260.

Brief overview of terms:

z196 and z114 package four processors on a single chip. Each processor has a private Level-1 Instruction cache, a private Level-1 Data cache and a private Level-2 cache that caches both instructions and data. These four processors on a chip share a Level-3 cache. When the Level-3 cache is described as On-Chip it physically resides on the same chip as the processor.

The z196 and z114 package six processor chips on a single book. z196 like z10 can contain up to four physical books in a system. z114 can contain up to two. Each book contains a Level-4 cache. When the Level-3 cache is described as Off Chip/On Book the cache is on a separate chip from the processor but on the same book. When a Level-4 cache is described as On Book it is on the same book as the processor, Off Book means the Level-4 cache is on a different book than the processor. Local memory is memory that is physically attached to the Book the processor is on. Remote memory is memory that is physically attached to a Book different from the processor.

Counter 128 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from the Level-2 cache.

Counter 129 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from the Level-2 cache.

Counter 130 – Level-1 Data TLB miss in progress. Incremented by one for every cycle a DTLB1 miss is in progress.

Counter 131 – Level-1 Instruction TLB miss in progress. Incremented by one for every cycle an ITLB1 miss is in progress.

Counter 132 – Undefined.

Counter 133 – Incremented by one for every store sent to Level-2 cache.

Counter 134 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off Book Level-3 cache.

Counter 135 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On Book Level-4 cache.

Counter 136 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On Book Level-4 cache.

Counter 137 – A directory write to the Level-1 Data cache where the line was originally in a Read-Only state in the cache but has been updated to be in the Exclusive state that allows stores to the cache line.

Counter 138 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off Book Level-4 cache.

Counter 139 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off Book Level-4 cache.

Counter 140 – A translation entry has been written to the Level-1 Data Translation Lookaside Buffer for a one-megabyte page.

Counter 141 – A directory write to the Level-1 Data cache where the installed cache line was sourced from memory that is attached to the same book as the Data cache (Local Memory).

Counter 142 – A directory write to the Level-1 Instruction cache where the installed cache line was sourced from memory that is attached to the same book as the Instruction cache (Local Memory).

Counter 143 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off Book Level-3 cache.

Counter 144 – A translation entry has been written to the Level-1 Data Translation Lookaside Buffer (DTLB1).

Counter 145 – A translation entry has been written to the Level-1 Instruction Translation Lookaside Buffer (ITLB1).

Counter 146 – A translation entry has been written to the Level-2 TLB Page Table Entry arrays.

Counter 147 – A translation entry has been written to the Level-2 TLB Combined Region Segment Table Entry arrays for a one-megabyte large page translation.

Counter 148 – A translation entry has been written to the Level-2 TLB Combined Region Segment Table Entry arrays.

Counter 149 – Undefined.

Counter 150 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Chip Level-3 cache.

Counter 151 – Undefined.

Counter 152 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off Chip/On Book Level-3 cache.

Counter 153 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Chip Level-3 cache.

Counter 154 – Undefined.

Counter 155 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off Chip/On Book Level-3 cache.

Counter 156 – Undefined.

To derive approximate rates for remote memory on Systems with more than one Book:

Directory writes to the Level-1 Instruction cache that was sourced from Remote Memory =

Counter 2 – (Counter 129 + 136 + 139 + 142 + 143 + 153 + 155)

Directory writes to the Level-1 Data cache that was sourced from Remote Memory =

Counter 4 – (Counter 128 + 134 + 135 + 138 + 141 + 150 + 152)

Extended Counters Definitions for zEC12 and zBC12

The Extended Counters provide information about hardware facilities and structures that are specific to a machine family. While much of the information provided in the Extended Counters for zEC12 are similar in kind to z10, z196/z114 details of the machines' micro-architectures differ and the zEC12 Counter definitions are not the same as z10, z196/z114.

zEC12 and zBC12 have the same counter definitions.

The Extended Counters are designed to expand upon information provided by the Basic Counter Set that is described in the Set-Program-Parameter and CPU-Measurement Facilities, SA23-2260.

Brief overview of terms:

zEC12 can package six processors on a single chip. Each processor has a private Level-1 Instruction cache, a private Level-1 Data cache, a private Level-2 Instruction cache and a private Level-2 Data cache. These six processors on a chip share a Level-3 cache. When the Level-3 cache is described as On-Chip it physically resides on the same chip as the processor.

The zEC12 can package six processor chips on a single book. zEC12, like z196 and z10 can contain up to four physical books in a system. Each book contains a Level-4 cache. When the Level-3 cache is described as Off Chip/On Book the cache is on a separate chip from the processor but on the same book. When a Level-4 cache is described as On Book it is on the same book as the processor, Off Book means the Level-4 cache is on a different book than the processor. Local memory is memory that is physically attached to the Book the processor is on. Remote memory is memory that is physically attached to a Book different from the processor.

An intervention indicates that another processor in the system was holding a cache line exclusive for stores and is a normal part of cache operation. z10, z196/z114 did not break these cases out separately in their extended counters, zEC12 does.

Counter 128 – Level-1 Data TLB miss in progress. Incremented by one for every cycle a DTLB1 miss is in progress.

Counter 129 – Level-1 Instruction TLB miss in progress. Incremented by one for every cycle an ITLB1 miss is in progress.

Counter 130 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from the Level-2 Instruction cache.

Counter 131 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from the Level-2 Instruction cache.

Counter 132 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from the Level-2 Data cache.

Counter 133 – A translation entry has been written to the Level-1 Data Translation Lookaside Buffer (DTLB1).

Counter 134 – Undefined.

Counter 135 - A directory write to the Level-1 Data cache where the installed cache line was sourced from memory that is attached to the same book as the Data cache (Local Memory).

Counter 136 - Undefined

Counter 137 - A directory write to the Level-1 Instruction cache where the installed cache line was sourced from memory that is attached to the same book as the Instruction cache (Local Memory).

Counter 138 – A directory write to the Level-1 Data cache where the line was originally in a Read-Only state in the cache but has been updated to be in the Exclusive state that allows stores to the cache line.

Counter 139 – A translation entry has been written to the Level-1 Data Translation Lookaside Buffer for a one-megabyte page.

Counter 140 – A translation entry has been written to the Level-1 Instruction Translation Lookaside Buffer (ITLB1).

Counter 141 – A translation entry has been written to the Level-2 TLB Page Table Entry arrays.

Counter 142 – A translation entry has been written to the Level-2 TLB Combined Region Segment Table Entry arrays for a one-megabyte large page translation.

Counter 143 – A translation entry has been written to the Level-2 TLB Combined Region Segment Table Entry arrays.

Counter 144 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Chip Level-3 cache without intervention.

Counter 145 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off Chip/On Book Level-3 cache without intervention.

Counter 146 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off Book Level-3 cache without intervention.

Counter 147 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On Book Level-4 cache.

Counter 148 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off Book Level-4 cache.

Counter 149 – A TEND instruction has completed in a nonconstrained transactional-execution mode.

Counter 150 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Chip Level-3 cache with intervention.

Counter 151 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off Chip/On Book Level-3 cache with intervention.

Counter 152 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off Book Level-3 cache with intervention.

Counter 153 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Chip Level-3 cache without intervention.

Counter 154 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off Chip/On Book Level-3 cache without intervention.

Counter 155 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off Book Level-3 cache without intervention.

Counter 156 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On Book Level-4 cache.

Counter 157 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off Book Level-4 cache.

Counter 158 – A TEND instruction has completed in a constrained transactional-execution mode.

Counter 159 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Chip Level-3 cache with intervention.

Counter 160 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off Chip/On Book Level-3 cache with intervention.

Counter 161 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off Book Level-3 cache with intervention.

Counter 162 to 176 – Undefined

Counter 177 – A transaction abort has occurred in a nonconstrained transactional-execution mode.

Counter 178 – A transaction abort has occurred in a constrained transactional-execution mode and the CPU is not using any special logic to allow the transaction to complete.

Counter 179 - A transaction abort has occurred in a constrained transactional-execution mode and the CPU is using special logic to allow the transaction to complete.

Counter 180 to 183 – Undefined

The TLB that caches 2 gigabyte page translations does not have an extended counter.

To derive approximate rates for remote memory for Systems with more than one Book:

Directory writes to the Level-1 Instruction cache that was sourced from Remote Memory =

Counter 2 – (Counter 131 + 137 + 153 + 154 + 155 + 156 + 157 + 159 + 160 + 161)

Directory writes to the Level-1 Data cache that was sourced from Remote Memory =

Counter 4 – (Counter 130 + 132 + 135 + 144 + 145 + 146 + 147 + 148 + 150 + 151 + 152)

Extended Counters for z13 and z13s

The Extended Counters provide information about hardware facilities and structures that are specific to a machine family. While much of the information provided in the Extended Counters for z13 are similar in kind to z10, z196/z114, zEC12 details of the machines' micro-architectures differ and the z13 Counter definitions are not the same as z10, z196/z114, zEC12.

The Extended Counters are designed to expand upon information provided by the Basic Counter Set that is described in the Set-Program-Parameter and CPU-Measurement Facilities, SA23-2260.

Brief overview of terms:

z13 provides Multi-threading. With multi-threading each processor (Core) on a chip can provide two logical processors or threads, (CPUs) for certain engine types. The Extended counters are provided on a per thread or logical CPU basis. Because of this the term CPU is used for z13 while the term processor was used for z10, z196/z114, zEC12/zBC12.

z13 can package eight Cores on a single chip, each of which can have two logical CPUs. Each Core has a private Level-1 Instruction cache, a private Level-1 Data cache, a private Level-2 Instruction cache and a private Level-2 Data cache.

These eight Cores on a chip share a Level-3 cache. When the Level-3 cache is described as On-Chip it physically resides on the same chip as the CPU.

z13 can package three Core chips on a single node. Each node contains a Level-4 cache. When a Level-3 or Level-4 cache is described as On-Node it physically resides on the same node as the CPU.

Each node can be physically packaged together with another node to make up a physical drawer. When a Level-3 or Level-4 cache is described as On-Drawer it physically resides on the same drawer as the CPU, but is not contained within the same physical node.

z13 can contain up to four drawers in a system. Connections between drawers are separated into

Same-Column and Far-Column connections. When a Level-3 cache is described as Off-Drawer Same-Column it is getting its data from a different drawer than the CPU is on and using a Same-Column connection. When a Level-3 cache is described as Off-Drawer Far-Column it is getting its data from a different drawer using a Far-Column connection.

An intervention indicates that another processor in the system was holding a cache line exclusive for stores and is a normal part of cache operation. The extended counters included these cases for zEC12/zBC12 and z13 continues to do so.

z13 provides more direct connectivity than prior systems. Memory can be directly attached to the chip the CPU resides on. It can be in the same node the CPU is on. It can also be on the same drawer or a different drawer. When memory is described as On-Chip, it is physically connected to the chip the CPU is on. On-Node memory is on the same physical node as the CPU. On-Drawer memory is on the same physical drawer as the CPU. Off-Drawer the memory is physically attached to a different drawer than the CPU is on.

Counter 128 – A directory write to the Level-1 Data cache where the line was originally in a Read-Only state in the cache but has been updated to be in the Exclusive state that allows stores to the cache line.

Counter 129 – A translation entry has been written to a Level-1 Data Translation Lookaside Buffer (DTLB1).

Counter 130 – Level-1 Data TLB miss in progress. Incremented by one for every cycle a DTLB1 miss is in progress.

Counter 131 – A translation entry has been written to the Level-1 Data Translation Lookaside Buffer for a one-megabyte page.

Counter 132 – A translation entry has been written to the Level-1 Data Translation Lookaside Buffer for a two-gigabyte page.

Counter 133 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from the Level-2 Data cache.

Counter 134 – A translation entry has been written to the Level-1 Instruction Translation Lookaside Buffer (ITLB1).

Counter 135 – Level-1 Instruction TLB miss in progress. Incremented by one for every cycle an ITLB1 miss is in progress.

Counter 136 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from the Level-2 Instruction cache.

Counter 137 – A translation entry has been written to the Level-2 TLB Page Table Entry arrays.

Counter 138 – A translation entry has been written to the Level-2 TLB Combined Region Segment Table Entry arrays for a one-megabyte large page translation.

Counter 139 – A translation entry has been written to the Level-2 TLB Combined Region Segment Table Entry arrays.

Counter 140 – A TEND instruction has completed in a constrained transactional-execution mode.

Counter 141 – A TEND instruction has completed in a non-constrained transactional-execution mode.

Counter 142 – Undefined

Counter 143 – Increments by one for any cycle where a level-1 cache or level-1 TLB miss is in progress.

Counter 144 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Chip Level-3 cache without intervention.

Counter 145 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Chip Level-3 cache with intervention.

Counter 146 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Node Level-4 cache.

Counter 147 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Node Level-3 cache with intervention.

Counter 148 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Node Level-3 cache without intervention.

Counter 149 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Drawer Level-4 cache.

Counter 150 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Drawer Level-3 cache with intervention.

Counter 151 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Drawer Level-3 cache without intervention.

Counter 152 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off-Drawer Same-Column Level-4 cache.

Counter 153 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off-Drawer Same-Column Level-3 cache with intervention.

Counter 154 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off-Drawer Same-Column Level-3 cache without intervention.

Counter 155 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off-Drawer Far-Column Level-4 cache.

Counter 156 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off-Drawer Far-Column Level-3 cache with intervention.

Counter 157 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off-Drawer Far-Column Level-3 cache without intervention.

Counter 158 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from On-Node memory.

Counter 159 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from On-Drawer memory.

Counter 160 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from Off-Drawer memory.

Counter 161 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from On-Chip memory.

Counter 162 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Chip Level-3 cache without intervention.

Counter 163 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Chip Level-3 cache with intervention.

Counter 164 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Node Level-4 cache.

Counter 165 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Node Level-3 cache with intervention.

Counter 166 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Node Level-3 cache without intervention.

Counter 167 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Drawer Level-4 cache.

Counter 168 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Drawer Level-3 cache with intervention.

Counter 169 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Drawer Level-3 cache without intervention.

Counter 170 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off-Drawer Same-Column Level-4 cache.

Counter 171 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off-Drawer Same-Column Level-3 cache with intervention.

Counter 172 - A directory write to the Level-1 Instruction cache directory where the returned cache

line was sourced from an Off-Drawer Same-Column Level-3 cache without intervention.

Counter 173 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off-Drawer Far-Column Level-4 cache.

Counter 174 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off-Drawer Far-Column Level-3 cache with intervention.

Counter 175 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off-Drawer Far-Column Level-3 cache without intervention.

Counter 176 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from On-Node memory.

Counter 177 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from On-Drawer memory.

Counter 178 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from Off-Drawer memory.

Counter 179 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from On-Chip memory.

Counters 180 to 217 - Undefined

Counter 218 - A transaction abort has occurred in a non-constrained transactional-execution mode.

Counter 219 - A transaction abort has occurred in a constrained transactional-execution mode and the CPU is not using any special logic to allow the transaction to complete.

Counter 220 - A transaction abort has occurred in a constrained transactional-execution mode and the CPU is using special logic to allow the transaction to complete.

Counter 221 to 225 - Undefined.

In addition to Extended Counters there is the MT-Diagnostic Counter set described in SA22-2260.

Counter 448 – Cycle count with one thread active.

Counter 449 – Cycle count with two threads active.

Counter 450-495 Undefined.

To derive approximate rates for remote memory for Systems with more than one Book:

Directory writes to the Level-1 Instruction cache that was sourced from Remote Memory =

Counter 177 + 178

Directory writes to the Level-1 Data cache that was sourced from Remote Memory =

Counter 159 + 160

Extended Counters for z14

The Extended Counters provide information about hardware facilities and structures that are specific to a machine family. While much of the information provided in the Extended Counters for z14 are similar in kind to z10, z196/z114, zEC12/zBC12, z13/z13s details of the machines' micro-architectures differ and the z14 Counter definitions are not the same as z10, z196/z114, zEC12/zBC12, z13/z13s.

The Extended Counters are designed to expand upon information provided by the Basic Counter Set that is described in the Set-Program-Parameter and CPU-Measurement Facilities, SA23-2260.

Brief overview of terms:

z14 provides Multi-threading. With multi-threading each processor (Core) on a chip can provide two logical processors or threads, (CPUs) for certain engine types. The Extended counters are provided on a per thread or logical CPU basis. Because of this the term CPU is used for z14 while the term processor was used for z10, z196/z114, zEC12/zBC12.

z14 can package ten Cores on a single chip, each of which can have two logical CPUs. Each Core has a private Level-1 Instruction cache, a private Level-1 Data cache, a private Level-2 Instruction cache and a private Level-2 Data cache.

These ten Cores on a chip share a Level-3 cache. When the Level-3 cache is described as On-Chip it physically resides on the same chip as the CPU.

z14 packages three Core chips per Cluster. There is one Level-4 cache per two Clusters. When a Level-3 is described as On-Cluster it physically resides on the same Cluster as the CPU.

Two Clusters are physically packaged together to make up a physical drawer. When a Level-3 cache is described as Off-Cluster it is on the same physical drawer but not the same Cluster as the CPU.

When a Level-4 cache is described as On Drawer it physically resides on the same drawer as the CPU. z14 can contain up to four drawers in a system.

An intervention indicates that another processor in the system was holding a cache line exclusive for stores and is a normal part of cache operation. The extended counters included these cases for zEC12/zBC12/z13/z13s and z14 continues to do so.

z14 breaks out On-Chip Level-3 cache hits into two categories. The cache line is in the On-Chip Level-3 cache and the cache line is in the correct state for the operation to be done (exclusive for stores, read-only or exclusive for fetches) and the On-Chip Level-3 cache has the cache line but does not have it in the correct state (read-only state for a store operation). In this second case an invalidate is sent to other L3s that have read-only copies of the cache line. In z13 these cases were not broken out. In z14 you must add these two cases to equal what z13 counted as On-Chip Level-3 cache hits. The read-only invalidate case only happens for cache line installs in the Level-1 Data cache.

For z14 memory can be directly attached to the chip the CPU resides on. It can be in the same cluster the CPU is on. It can also be on the same drawer or a different drawer. When memory is described as On-Chip, it is physically connected to the chip the CPU is on. On-Cluster memory is on the same physical cluster as the CPU. On-Drawer memory is on the same physical drawer as the CPU but not the same cluster. Off-Drawer the memory is physically attached to a different drawer than the CPU is on.

On z14 the TLB/Cache structure has changed. What was the Level-1 TLBs on z13 have been merged into the Level-1 cache directories. The Level-2 TLB is accessed in tandem with Level-2 cache so that translation information and data can be transferred at the same time.

A Last Host Translation is the final step of the Level-2 TLB to get a Guest 1 Absolute address from a Guest 2 Real address on a Level-2 TLB miss. The Page Table Entry of the Last Host Translation is not written into the Level-2 TLB.

Counter 128 – A directory write to the Level-1 Data cache where the line was originally in a Read-Only state in the cache but has been updated to be in the Exclusive state that allows stores to the cache line.

Counter 129 – A translation has been written into The Translation Lookaside Buffer 2 (TLB2) and the request was made by the data cache. This is a replacement for what was provided for the DTLB on prior machines.

Counter 130 – A TLB2 miss is in progress for a request made by the data cache. Incremented by one for every TLB2 miss in progress for the Level-1 Data cache on this cycle. This is a replacement for what was provided for the DTLB on prior machines.

Counter 131 – A translation entry was written into the Combined Region and Segment Table Entry array in the Level-2 TLB for a one-megabyte page or a Last Host Translation was done.

Counter 132 – A translation entry for a two-gigabyte page was written into the Level-2 TLB.

Counter 133 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from the Level-2 Data cache.

Counter 134 – A translation entry has been written into the Translation Lookaside Buffer 2 (TLB2) and the request was made by the instruction cache. This is a replacement for what was provided for the ITLB on prior machines.

Counter 135 – A TLB2 miss is in progress for a request made by the instruction cache. Incremented by one for every TLB2 miss in progress for the Level-1 Instruction cache in a cycle. This is a replacement for what was provided for the ITLB on prior machines.

Counter 136 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from the Level-2 Instruction cache.

Counter 137 – A translation entry was written into the Page Table Entry array in the Level-2 TLB.

Counter 138 – Translation entries were written into the Combined Region and Segment Table Entry array and the Page Table Entry array in the Level-2 TLB.

Counter 139 – The number of Level-2 TLB translation engines busy in a cycle.

Counter 140 – A TEND instruction has completed in a constrained transactional-execution mode.

Counter 141 – A TEND instruction has completed in a non-constrained transactional-execution mode.

Counter 142 – Undefined

Counter 143 – Increments by one for any cycle where a level-1 cache or level-2 TLB miss is in progress.

Counter 144 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Chip Level-3 cache without intervention.

Counter 145 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from On-Chip memory.

Counter 146 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Chip Level-3 cache with intervention.

Counter 147 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from On-Cluster Level-3 cache without intervention.

Counter 148 - A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Cluster memory.

Counter 149 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an On-Cluster Level-3 cache with intervention.

Counter 150 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off-Cluster Level-3 cache without intervention.

Counter 151 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from Off-Cluster memory.

Counter 152 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off-Cluster Level-3 cache with intervention.

Counter 153 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off-Drawer Level-3 cache without intervention.

Counter 154 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from Off-Drawer memory.

Counter 155 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from an Off-Drawer Level-3 cache with intervention.

Counter 156 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from On-Drawer Level-4 cache.

Counter 157 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from Off-Drawer Level-4 cache.

Counter 158 – A directory write to the Level-1 Data cache directory where the returned cache line was sourced from On-Chip L3 but a read-only invalidate was done to remove other copies of the cache line.

Counter 159 to 161 – Undefined

Counter 162 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Chip Level-3 cache without intervention.

Counter 163 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from On-Chip memory.

Counter 164 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Chip Level-3 cache with intervention.

Counter 165 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Cluster Level-3 cache without intervention.

Counter 166 - A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an On-Cluster memory.

Counter 167 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from On-Cluster Level-3 cache with intervention.

Counter 168 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off-Cluster Level-3 cache without intervention.

Counter 169 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from Off-Cluster memory.

Counter 170 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off-Cluster Level-3 cache with intervention.

Counter 171 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off-Drawer Level-3 cache without intervention.

Counter 172 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from Off-Drawer memory.

Counter 173 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from an Off-Drawer Level-3 cache with intervention.

Counter 174 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from On-Drawer Level-4 cache.

Counter 175 – A directory write to the Level-1 Instruction cache directory where the returned cache line was sourced from Off-Drawer Level-4 cache.

Counter 176 to 223 – Undefined

Counter 224 – Count of floating point execution slots used for finished Binary Coded Decimal to Decimal Floating Point conversions. Instructions: CDZT, CXZT, CZDT, CZXT.

Counter 225 – Count of floating point execution slots used for finished vector arithmetic Binary Coded Decimal instructions. Instructions: VAP, VSP, VMP, VMSP, VDP, VSDP, VRP, VLIP, VSRP, VPSOP, VCP, VTP, VPKZ, VUPKZ, VCVB, VCVBG, VCVD, VCVDG.

Counter 226 – Decimal instructions dispatched. Instructions: CVB, CVD, AP, CP, DP, ED, EDMK, MP, SRP, SP, ZAP.

Counters 227 to 231 – Undefined

Counters 232 – Last Host Translation done.

Counters 233 to 242 – Undefined

Counter 243 – A transaction abort has occurred in a non-constrained transactional-execution mode.

Counter 244 – A transaction abort has occurred in a constrained transactional-execution mode and the CPU is not using any special logic to allow the transaction to complete.

Counter 245 - A transaction abort has occurred in a constrained transactional-execution mode and the CPU is using special logic to allow the transaction to complete.

Counters 246 to 255 – Undefined

In addition to Extended Counters there is the MT-Diagnostic Counter set described in SA22-2260.

Counter 448 – Cycle count with one thread active.

Counter 449 – Cycle count with two threads active.

Counter 450-495 Undefined.

To replace Combined Region and Segment Table array writes from prior machines:

Counter 131 + Counter 138

To replace Combined Region and Segment Table array writes for one-megabyte page translations on prior machines:

Counter 131 – Counter 232

Depending on the mix of translations at Guest 2 the above calculation may under count writes for one-megabyte page translations.

To replace Page Table Entry array writes on prior machines:

Counter 137 + Counter 138

Extended Counters for z15

The Extended Counters provide information about hardware facilities and structures that are specific to a machine family.

The Extended Counters are designed to expand upon information provided by the Basic Counter Set that is described in the Set-Program-Parameter and CPU-Measurement Facilities, SA23-2260.

Brief overview of terms:

z15 provides Multi-threading. With multi-threading each processor (Core) on a chip can provide two logical processors or threads, (CPUs) for certain engine types. The Extended counters are provided on a per thread or logical CPU basis. Because of this the term CPU is used for z15 while the term processor was used for z10, z196/z114, zEC12/zBC12.

z15 can package twelve Cores on a single chip, each of which can have two logical CPUs. Each Core has a private Level-1 Instruction cache, a private Level-1 Data cache, a private Level-2 Instruction cache and a private Level-2 Data cache.

These twelve Cores on a chip share a Level-3 cache. When the Level-3 cache is described as On-Chip it physically resides on the same chip as the CPU.

z15 packages two Core chips per Cluster. There is one Level-4 cache per two Clusters. When a Level-3 cache is described as On-Cluster it physically resides on the same Cluster as the CPU.

Two Clusters are physically packaged together to make up a physical drawer. When a Level-3 cache is described as Off-Cluster it is on the same physical drawer but not the same Cluster as the CPU.

When a Level-4 cache is described as On-Drawer it physically resides on the same drawer as the CPU. z15 can contain up to five drawers in a system.

An intervention indicates that another processor in the system was holding a cache line exclusive for stores and is a normal part of cache operation. The extended counters included these cases for zEC12/zBC12/z13/z13s/z14 and z15 continues to do so.

z15 breaks out On-Chip Level-3 cache hits into two categories. The cache line is in the On-Chip Level-3 cache and the cache line is in the correct state for the operation to be done (exclusive for stores, read-only or exclusive for fetches) and the On-Chip Level-3 cache has the cache line but does not have it in the correct state (read-only state for a store operation). In this second case an invalidate is sent to other L3s that have read-only copies of the cache line. In z13 these cases were not broken out. In z14 and z15 you must add these two cases to equal what z13 counted as On-Chip Level-3 cache hits. The read-only invalidate case only happens for cache line installs in the Level-1 Data cache.

For z15 memory can be directly attached to the chip the CPU resides on. It can be in the same cluster the CPU is on. It can also be on the same drawer or a different drawer. When memory is described as On-Chip, it is physically connected to the chip the CPU is on. On-Cluster memory is on the same physical cluster as the CPU. On-Drawer memory is on the same physical drawer as the CPU but not the same cluster. Off-Drawer the memory is physically attached to a different drawer than the CPU is on.

On z15 the TLB/Cache structure is the same as z14. What was the Level-1 TLBs on z13 have been merged into the Level-1 cache directories. The Level-2 TLB is accessed in tandem with Level-2 cache so that translation information and data can be transferred at the same time.

A Last Host Translation is the final step of the Level-2 TLB to get a Guest 1 Absolute address from a Guest 2 Real address on a Level-2 TLB miss. The Page Table Entry of the Last Host Translation is not written into the Level-2 TLB.

For all counters not described in the z15 section of this document refer to the z14 section for definition.

Counter 131 – A translation entry was written into the Combined Region and Segment Table Entry array in the Level-2 TLB for a one-megabyte page.

Counter 246 – Undefined

Counter 247 – cycles CPU spent obtaining access to Deflate unit

Counters 248 to 251 – Undefined

Counter 252 – cycles CPU is using Deflate unit

Counters 253 to 263 – Undefined

Counter 264 – Increments by one for every DEFLATE CONVERSION CALL instruction executed.

Counter 265 – Increments by one for every DEFLATE CONVERSION CALL instruction executed that ended in Condition Codes 0, 1 or 2.

Counters 266 to 271 – Undefined

To replace Combined Region and Segment Table array writes for one-megabyte page translations on z13 and prior machines:

Counter 131

End of Document

SA23-2261-06

