

CPU MF Formulas and Updates

December 2023

John Burg



z/OS SMF 113 Record

- SMF113_2_CTRVN2
 - “1” = z10
 - “2” = z196 / z114
 - “3” = zEC12 / zBC12
 - “4” = z13 / z13s
 - “5” = z14
 - “6” = z15
 - “7” = z16

L1MP and RNI-based LSPR Workload Decision Table

L1MP	RNI	LSPR Workload Match
< 3%	≥ 0.75	AVERAGE
	< 0.75	LOW
3% to 6%	> 1.0	HIGH
	0.6 to 1.0	AVERAGE
	< 0.6	LOW
> 6%	≥ 0.75	HIGH
	< 0.75	AVERAGE

Current table applies to z10 EC, z10 BC, z196, z114, zEC12, zBC12, z13, z13s, z14, z15 and z16 CPU MF data

z16 Metrics

z16 vs z15 Hardware and Topology Comparison


z15 CPU 5.2 GHz
 Caches L1 private 128k i, 128k d / core
 L2 private 4 MB i, 4 MB d / core
 L3 shared 256 MB / CP chip
 L4 shared 960 MB / drawer

z16 CPU 5.2 GHz
 Caches L1 private 128k i, 128k d / core
 L2 private 32 MB unified / core
 virtual victim L3 up to 7x32 = 224 MB / CP chip
 virtual victim L4 up to 8x32x7 = 1.75 GB / drawer


Topology
 - 12 cores + 1 L3 / CP chip
 - 2 CP chips / cluster
 - 2 clusters + 1 L4 (48 engines) / drawer
 - 5 drawers / CEC
 - Book interconnect: numa star

Topology
 - 8 (core + L3)s / CP chip
 - 2 CP chips / DCM
 - 4 DCMs (64 engines) / drawer
 - 4 drawers / CEC
 - Book interconnect: numa star


z15 Single Drawer View (physical+logical)



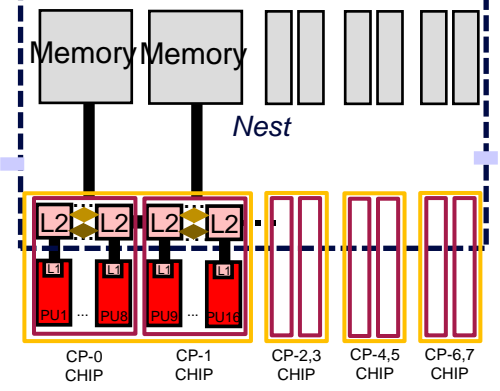
z16 Single CP Chip View (physical)




z16 Single Drawer View (physical)



z16 Single Drawer View (physical)



z16 Single Drawer View (logical)



Formulas – z16

Workload Characterization
L1 Sourcing from cache/memory hierarchy

Metric	Calculation – note all fields are <i>deltas</i> . SMF113-1s are deltas. SMF 113-2s are cumulative.
CPI	$B0 / B1$
PRBSTATE	$(P33 / B1) * 100$
L1MP	$((B2+B4) / B1) * 100$
L2P	$((E145+E146+E169+E170) / (B2+B4)) * 100$
L3P	$((E147+E149+E150+E151+E171+E173+E174+E175) / (B2+B4)) * 100$
L4LP	$((E148+E152+E153+E154+E160+E161+E162+E163+E164+E165+E172+E176+E177+E178) / (B2+B4)) * 100$
L4RP	$((E155+E166+E167+E168+E179) / (B2+B4)) * 100$
MEMP	$((E156+E157+E158+E159+E180+E181+E182+E183) / (B2+B4)) * 100$
LPARCPU	$(((1/CPSP/1,000,000) * B0) / Interval in Seconds) * 100$

Updated May 31, 2022

Note these Formulas may change in the future

CPI – Cycles per Instruction
Prb State - % Problem State
L1MP – Level 1 Miss Per 100 instructions
L2P – % sourced from Level 2 cache
L3P – % sourced from Level 3 on same Chip cache
L4LP – % sourced from Level 4 Local cache (on same book)
L4RP – % sourced from Level 4 Remote cache (on different book)
MEMP - % sourced from Memory
LPARCPU - APPL% (GCPs, ZAAPs, zIIPs) captured and uncaptured

B* - Basic Counter Set - Counter Number
P* - Problem-State Counter Set - Counter Number
See "The Load-Program-Parameter and CPU-Measurement Facilities" SA23-2260 for full description
E* - Extended Counters - Counter Number
See "IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/ z114, zEC12/zBC12, z13/z13s, z14, z15 and z16" SA23-2261-07 for full description
CPSP - SMF113_2_CPSP "CPU Speed"

Formulas – z16 Additional

Metric	Calculation- <i>note all fields are <u>deltas</u>. SMF113-1s are deltas. SMF 113-2s are cumulative.</i>
Est Instr Cmplx CPI	CPI – Finite CPI
Finite CPI	E143 / B1
SCPL1M	E143 / (B2+B4)
Rel Nest Intensity	$4.3 * (0.45 * L3P + 1.3 * L4LP + 5.0 * L4RP + 6.1 * MEMP) / 100$
Eff GHz	CPSP / 1000

Updated May 31, 2022

Note these Formulas may change in the future

Est Instr Cmplx CPI – Estimated Instruction Complexity CPI (infinite L1)

Est Finite CPI – Estimated CPI from Finite cache/memory

Est SCPL1M – Estimated Sourcing Cycles per Level 1 Miss

Rel Nest Intensity – Reflects distribution and latency of sourcing from shared caches and memory

Eff GHz – Effective gigahertz for GCPs, cycles per nanosecond

B* - Basic Counter Set - Counter Number


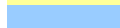
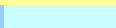
P* - Problem-State Counter Set - Counter Number

See "The Load-Program-Parameter and CPU-Measurement Facilities" SA23-2260 for full description

E* - Extended Counters - Counter Number

See "IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/z114, zEC12/zBC12, z13/z13s, z14, z15 and z16" SA23-2261-07 for full description

CPSP - SMF113_2_CPSP "CPU Speed"

 Workload Characterization
  L1 Sourcing from cache/memory hierarchy

Formulas – z16 Additional TLB

Metric	Calculation – note all fields are <i>deltas</i> . SMF113-1s are deltas. SMF 113-2s are cumulative.
Est. TLB1 CPU Miss % of Total CPU	$((E130+E135) / B0) * (E143 / (B3+B5)) * 100$
Estimated TLB1 Cycles per TLB Miss	$(E130+E135) / (E129+E134) * (E143 / (B3+B5))$
PTE % of all TLB1 Misses	N/A with processor design change
TLB Miss Rate	$(E129 + E134) / \text{interval}$

Updated May 31, 2022

Note these Formulas may change in the future

Est. TLB1 CPU Miss % of Total CPU - Estimated TLB CPU % of Total CPU

Estimated TLB1 Cycles per TLB Miss – Estimated Cycles per TLB Miss

PTE % of all TLB1 Misses – Page Table Entry % misses

TLB Miss Rate – TLB Misses per interval (interval is defined by user for length of measurement and units)

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See "The Load-Program-Parameter and CPU-Measurement Facilities" SA23-2260 for full description

E* - Extended Counters - Counter Number

See "IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/z114, zEC12 /zBC12, z13/z13s, z14, z15 and z16" SA23-2261-07 for full description

CPSP - SMF113_2_CPSP "CPU Speed"

Formulas – z16 Additional AIU

Metric	Calculation – note all fields are <u>deltas</u> . SMF113-1s are deltas. SMF 113-2s are cumulative.
W_AIU_CPU	$((1/CPSP/1,000,000) * E269) / \text{Interval in Seconds} * 100$
C_AIU_CPU	$((1/CPSP/1,000,000) * E270) / \text{Interval in Seconds} * 100$
AIU_CPU	W_AIU_CPU + C_AIU_CPU

Updated Aug 14, 2023

Note these Formulas may change in the future

Counter 267 – Increments by one for every NUAL NETWORK PROCESSING ASSIST instruction executed.

Counter 268 – Increments by one for every NUAL NETWORK PROCESSING ASSIST instruction executed that ended in Condition Codes 0, 1 or 2.

Counter 269 – Cycles CPU spent obtaining access to IBM Z Integrated Accelerator for AI.

Counter 270 – Cycles CPU is using IBM Z Integrated Accelerator for AI

W_AIU_CPU - Waiting for access to AIU (LPARCPU units)

C_AIU_CPU – Executing AIU (LPARCPU units)

AIU_CPU – Total AIU CPU (LPARCPU units)

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See "The Load-Program-Parameter and CPU-Measurement Facilities" SA23-2260 for full description

E* - Extended Counters - Counter Number

See "IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/z114, zEC12/zBC12, z13/z13s, z14, z15 and z16" SA23-2261-07 for full description

CPSP - SMF113_2_CPSP "CPU Speed"

z15 Metrics

z14 vs z15 Hardware Comparison

■ z14 (3906)

- CPU (14nm SOI)
 - 5.2 GHz
- Caches
 - L1 private 128k i, 128k d
 - L2 private 2 MB i, 4 MB d
 - L3 shared 128 MB per chip
 - L4 shared 672 MB per drawer

■ Topology

- 10 cores + 1 L3 per CP chip
- 2-or-3 CP chips per cluster
- 2 clusters + 1 L4 per drawer
- 4 drawers max per CPC
- Book interconnect: NUMA star


■ z15 (8561)

- CPU (14 nm SOI)
 - 5.2 GHz
- Caches
 - L1 private 128k i, 128k d
 - L2 private **4 MB i**, 4 MB d
 - L3 shared **256 MB** per chip
 - L4 shared **960 MB** per drawer

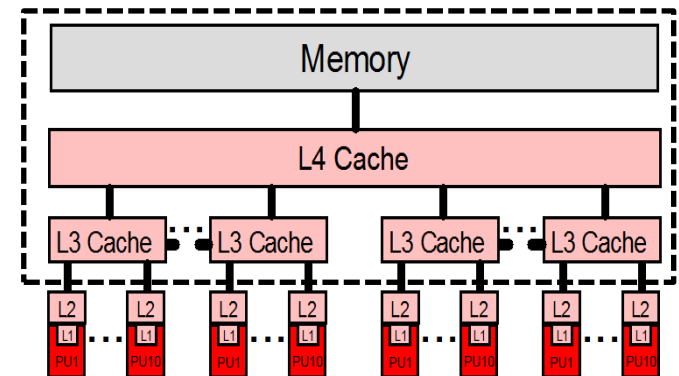
■ Topology

- **12** cores + 1 L3 per CP chip
- **2** CP chips per cluster
- 2 clusters + 1 L4 per drawer
- **5** drawers max per CPC
- Book interconnect: NUMA star

Single Drawer View



Single Drawer View



Formulas – z15

Workload Characterization
 L1 Sourcing from cache/memory hierarchy

Metric	Calculation – note all fields are <u>deltas</u> . SMF113-1s are deltas. SMF 113-2s are cumulative.
CPI	$B0 / B1$
PRBSTATE	$(P33 / B1) * 100$
L1MP	$((B2+B4) / B1) * 100$
L2P	$((E133+E136) / (B2+B4)) * 100$
L3P	$((E144+E146+E162+E164) / (B2+B4)) * 100$
L4LP	$((E147+E149+E156+E165+E167+E174+E150+E152+E158+E168+E170) / (B2+B4)) * 100$
L4RP	$((E153+E155+E157+E171+E173+E175) / (B2+B4)) * 100$
MEMP	$(((E145 + E148 + E151 + E154 + E163 + E166 + E169 + E172) / (B2+B4)) * 100$
LPARCPU	$(((1/CPSP/1,000,000) * B0) / Interval \textit{ in Seconds}) * 100$

CPI – Cycles per Instruction
 Prb State - % Problem State
 L1MP – Level 1 Miss Per 100 instructions
 L2P – % sourced from Level 2 cache
 L3P – % sourced from Level 3 on same Chip cache
 L4LP – % sourced from Level 4 Local cache (on same book)
 L4RP – % sourced from Level 4 Remote cache (on different book)
 MEMP - % sourced from Memory
 LPARCPU - APPL% (GCPs, ZAAPs, zIIPs) captured and uncaptured

B* - Basic Counter Set - Counter Number
 P* - Problem-State Counter Set - Counter Number
 See "The Load-Program-Parameter and CPU-Measurement Facilities" SA23-2260 for full description
 E* - Extended Counters - Counter Number
 See "IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/z114, zEC12/zBC12, z13/z13s, z14 and z15" SA23-2261-05 for full description
 CPSP - SMF113_2_CPSP "CPU Speed"

Updated September 23, 2019

Note these Formulas may change in the future

Formulas – z15 Additional

Metric	Calculation- <i>note all fields are <u>deltas</u>. SMF113-1s are deltas. SMF 113-2s are cumulative.</i>
Est Instr Cmplx CPI	CPI – Estimated Finite CPI
Est Finite CPI	$(E143 / B1) + .15$
Est SCPL1M	$\text{Estimated Finite CPI} / (L1MP / 100)$
Rel Nest Intensity	$2.9 * (0.45 * L3P + 1.5 * L4LP + 3.2 * L4RP + 6.5 * MEMP) / 100$
Eff GHz	CPSP / 1000

Updated December 7, 2023

Note these Formulas may change in the future

Est Instr Cmplx CPI – Estimated Instruction Complexity CPI (infinite L1)

Est Finite CPI – Estimated CPI from Finite cache/memory

Est SCPL1M – Estimated Sourcing Cycles per Level 1 Miss

Rel Nest Intensity – Reflects distribution and latency of sourcing from shared caches and memory

Eff GHz – Effective gigahertz for GCPs, cycles per nanosecond

B* - Basic Counter Set - Counter Number


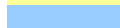
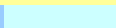
P* - Problem-State Counter Set - Counter Number

See "The Load-Program-Parameter and CPU-Measurement Facilities" SA23-2260 for full description

E* - Extended Counters - Counter Number

See "IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/ z114, zEC12 /zBC12, z13/z13s and z14" SA23-2261-04 for full description

CPSP - SMF113_2_CPSP "CPU Speed"

 Workload Characterization
  L1 Sourcing from cache/memory hierarchy

Formulas – z15 Additional TLB

Metric	Calculation – note all fields are <i>deltas</i> . SMF113-1s are <i>deltas</i> . SMF 113-2s are <i>cumulative</i> .
Est. TLB1 CPU Miss % of Total CPU	$((E130+E135) / B0) * (E143 / (B3+B5)) * 100$
Estimated TLB1 Cycles per TLB Miss	$(E130+E135) / (E129+E134) * (E143 / (B3+B5))$
PTE % of all TLB1 Misses	N/A with processor design change
TLB Miss Rate	$(E129 + E134) / \text{interval}$

Est. TLB1 CPU Miss % of Total CPU - Estimated TLB CPU % of Total CPU

Estimated TLB1 Cycles per TLB Miss – Estimated Cycles per TLB Miss

PTE % of all TLB1 Misses – Page Table Entry % misses

TLB Miss Rate – TLB Misses per interval (interval is defined by user for length of measurement and units)

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See "The Load-Program-Parameter and CPU-Measurement Facilities" SA23-2260 for full description

E* - Extended Counters - Counter Number

See "IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/ z114, zEC12 /zBC12, z13/z13s and z14" SA23-2261-04 for full description

CPSP - SMF113_2_CPSP "CPU Speed"

Updated September 23, 2019

Note these Formulas may change in the future

z14 Metrics

z14 vs z13 Hardware and Topology Comparison


z13

– CPU

- 5.0 GHz
- Major pipeline enhancements
- 1 picocoded translation engine

– Caches

- L1 private 96k i, 128k d
- L2 private 2 MB i, 2 MB d
- L3 shared 64 MB / chip
- L4 shared 480 MB / **node**
Plus 224 MB NIC




z14 – L3 clustering and cache sizes aside, topology strongly resembles zEC12

– CPU

- 5.2 GHz
- Logical directory w/ inclusive TLB
- 4 HW-implemented translation engines

– Caches

- L1 private 128k i, 128k d
- L2 private 2 MB i, 4 MB d
- L3 shared 128 MB / chip
- L4 shared 672 MB / **node drawer**



Formulas – z14

Workload Characterization
 L1 Sourcing from cache/memory hierarchy

Metric	Calculation – note all fields are <u>deltas</u> . SMF113-1s are deltas. SMF 113-2s are cumulative.
CPI	$B0 / B1$
PRBSTATE	$(P33 / B1) * 100$
L1MP	$((B2+B4) / B1) * 100$
L2P	$((E133+E136) / (B2+B4)) * 100$
L3P	$((E144+E146+E162+E164) / (B2+B4)) * 100$
L4LP	$((E147+E149+E156+E165+E167+E174+E150+E152+E158+E168+E170) / (B2+B4)) * 100$
L4RP	$((E153+E155+E157+E171+E173+E175) / (B2+B4)) * 100$
MEMP	$(((E145 + E148 + E151 + E154 + E163 + E166 + E169 + E172) / (B2+B4)) * 100$
LPARCPU	$(((1/CPSP/1,000,000) * B0) / Interval \textit{ in Seconds}) * 100$

CPI – Cycles per Instruction
 Prb State - % Problem State
 L1MP – Level 1 Miss Per 100 instructions
 L2P – % sourced from Level 2 cache
 L3P – % sourced from Level 3 on same Chip cache
 L4LP – % sourced from Level 4 Local cache (on same book)
 L4RP – % sourced from Level 4 Remote cache (on different book)
 MEMP - % sourced from Memory
 LPARCPU - APPL% (GCPs, ZAAPs, zIIPs) captured and uncaptured

B* - Basic Counter Set - Counter Number
 P* - Problem-State Counter Set - Counter Number
 See "The Load-Program-Parameter and CPU-Measurement Facilities" SA23-2260 for full description
 E* - Extended Counters - Counter Number
 See "IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/ z114, zEC12 /zBC12, z13/z13s and z14" SA23-2261-04 for full description
 CPSP - SMF113_2_CPSP "CPU Speed"

Updated December 2017

Note these Formulas may change in the future

Formulas – z14 Additional

Metric	Calculation– <i>note all fields are <u>deltas</u>. SMF113-1s are deltas. SMF 113-2s are cumulative.</i>
Est Instr Cmplx CPI	CPI – Estimated Finite CPI
Est Finite CPI	$(E143 / B1) + .18$
Est SCPL1M	$\text{Estimated Finite CPI} / (L1MP / 100)$
Rel Nest Intensity	$2.4 * (0.4 * L3P + 1.5 * L4LP + 3.2 * L4RP + 7.0 * MEMP) / 100$
Eff GHz	CPSP / 1000

Updated December 7, 2023

Note these Formulas may change in the future

Est Instr Cmplx CPI – Estimated Instruction Complexity CPI (infinite L1)

Est Finite CPI – Estimated CPI from Finite cache/memory

Est SCPL1M – Estimated Sourcing Cycles per Level 1 Miss

Rel Nest Intensity – Reflects distribution and latency of sourcing from shared caches and memory

Eff GHz – Effective gigahertz for GCPs, cycles per nanosecond

B* - Basic Counter Set - Counter Number


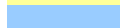
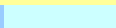
P* - Problem-State Counter Set - Counter Number

See "The Load-Program-Parameter and CPU-Measurement Facilities" SA23-2260 for full description

E* - Extended Counters - Counter Number

See "IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/ z114, zEC12 /zBC12, z13/z13s and z14" SA23-2261-04 for full description

CPSP - SMF113_2_CPSP "CPU Speed"

 Workload Characterization
  L1 Sourcing from cache/memory hierarchy

Formulas – z14 Additional TLB

Metric	Calculation – note all fields are <i>deltas</i> . SMF113-1s are <i>deltas</i> . SMF 113-2s are <i>cumulative</i> .
Est. TLB1 CPU Miss % of Total CPU	$((E130+E135) / B0) * (E143 / (B3+B5)) * 100$
Estimated TLB1 Cycles per TLB Miss	$(E130+E135) / (E129+E134) * (E143 / (B3+B5))$
PTE % of all TLB1 Misses	N/A with processor design change
TLB Miss Rate	$(E129 + E134) / \text{interval}$

Est. TLB1 CPU Miss % of Total CPU - Estimated TLB CPU % of Total CPU

Estimated TLB1 Cycles per TLB Miss – Estimated Cycles per TLB Miss

PTE % of all TLB1 Misses – Page Table Entry % misses

TLB Miss Rate – TLB Misses per interval (interval is defined by user for length of measurement and units)

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See "The Load-Program-Parameter and CPU-Measurement Facilities" SA23-2260 for full description

E* - Extended Counters - Counter Number

See "IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/ z114, zEC12 /zBC12, z13/z13s and z14" SA23-2261-04 for full description

CPSP - SMF113_2_CPSP "CPU Speed"

Updated September 23, 2019

Note these Formulas may change in the future

z13 Metrics

IBM z13 versus zEC12 Hardware Comparison

- zEC12


- CPU

- 5.5 GHz
 - Enhanced Out-Of-Order

- Caches

- L1 private 64k i, 96k d
 - L2 private 1 MB i + 1 MB d
 - L3 shared 48 MB / chip
 - L4 shared 384 MB / book

Single Book View



- z13


- CPU

- 5.0 GHz
 - Major pipeline enhancements

- Caches

- L1 private 96k i, 128k d
 - L2 private 2 MB i + 2 MB d
 - L3 shared 64 MB / chip
 - L4 shared 480 MB / node
 - plus 224 MB L3 NIC Directory

Single Drawer View - Two Nodes



Formulas – z13 / z13s

Workload Characterization
L1 Sourcing from cache/memory hierarchy

Metric	Calculation – note all fields are <i>deltas</i> . SMF113-1s are <i>deltas</i> . SMF 113-2s are <i>cumulative</i> .
CPI	$B0 / B1$
PRBSTATE	$(P33 / B1) * 100$
L1MP	$((B2+B4) / B1) * 100$
L2P	$((E133+E136) / (B2+B4)) * 100$
L3P	$((E144+E145+ E162+E163) / (B2+B4)) * 100$
L4LP	$((E146+E147+E148+E164+E165+E166) / (B2+B4)) * 100$
L4RP	$((E149+E150+E151+E152+E153+E154+E155+E156+E157+E167+E168+E169+E170+E171+E172+ E173+E174+E175) / (B2+B4)) * 100$
MEMP	$(((E158 + E159 + E160 + E161 + E176 + E177 + E178 + E179) / (B2+B4)) * 100$
LPARCPU	$(((1/CPSP/1,000,000) * B0) / Interval in Seconds) * 100$

CPI – Cycles per Instruction

Prb State - % Problem State

L1MP – Level 1 Miss Per 100 instructions

L2P – % sourced from Level 2 cache

L3P – % sourced from Level 3 on same Chip cache

L4LP – % sourced from Level 4 Local cache (on same book)

L4RP – % sourced from Level 4 Remote cache (on different book)

MEMP - % sourced from Memory

LPARCPU - APPL% (GCPs, zAAPs, zIIPs) captured and uncaptured

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See “The Load-Program-Parameter and CPU-Measurement Facilities” SA23-2260 for full description

E* - Extended Counters - Counter Number

See “IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/ z114, zEC12 /zBC12 and z13” SA23-2261-03 for full description

CPSP - SMF113_2_CPSP “CPU Speed”

Formulas – z13 / z13s Additional

Metric	Calculation- <i>note all fields are <u>deltas</u>. SMF113-1s are deltas. SMF 113-2s are cumulative.</i>
Est Instr Cmplx CPI	CPI – Estimated Finite CPI
Est Finite CPI	E143 / B1
Est SCPL1M	E143 / (B2+B4)
Rel Nest Intensity	2.3 *(0.4*L3P + 1.6*L4LP + 3.5*L4RP + 7.5*MEMP) / 100
Eff GHz	CPSP / 1000

Updated February 2017

Note these Formulas may change in the future


Est Instr Cmplx CPI – Estimated Instruction Complexity CPI (infinite L1)

Est Finite CPI – Estimated CPI from Finite cache/memory

Est SCPL1M – Estimated Sourcing Cycles per Level 1 Miss

Rel Nest Intensity –Reflects distribution and latency of sourcing from shared caches and memory

Eff GHz – Effective gigahertz for GCPs, cycles per nanosecond

 Workload Characterization
L1 Sourcing from cache/memory hierarchy

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See “The Load-Program-Parameter and CPU-Measurement Facilities” SA23-2260 for full description

E* - Extended Counters - Counter Number

See “IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/ z114, zEC12 /zBC12 and z13” SA23-2261-03 for full description

CPSP - SMF113_2_CPSP “CPU Speed”

Formulas – z13 / z13s Additional TLB

Metric	Calculation – note all fields are <u>deltas</u> . SMF113-1s are deltas. SMF 113-2s are cumulative.
Est. TLB1 CPU Miss % of Total CPU	$((E130+E135) / B0) * (E143 / (B3+B5)) * 100$
Estimated TLB1 Cycles per TLB Miss	$(E130+E135) / (E129+E134) * (E143 / (B3+B5))$
PTE % of all TLB1 Misses	$(E137 / (E129+E134)) * 100$
TLB Miss Rate	$(E129 + E134) / \text{interval}$

Note these Formulas may change in the future

Est. TLB1 CPU Miss % of Total CPU - Estimated TLB CPU % of Total CPU

Estimated TLB1 Cycles per TLB Miss – Estimated Cycles per TLB Miss

PTE % of all TLB1 Misses – Page Table Entry % misses

TLB Miss Rate – TLB Misses per interval (interval is defined by user for length of measurement and units)

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See “The Load-Program-Parameter and CPU-Measurement Facilities” SA23-2260 for full description

E* - Extended Counters - Counter Number

See “IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196/ z114, zEC12 /zBC12 and z13” SA23-2261-03 for full description

CPSP - SMF113_2_CPSP “CPU Speed”

zEC12 Metrics

zEC12 versus z196 hardware comparison


■ z196

▶ CPU

- 5.2 GHz
- Out-Of-Order execution

▶ Caches

- L1 private 64k i, 128k d
- L2 private 1.5 MB
- L3 shared 24 MB / chip
- L4 shared 192 MB / book




■ zEC12

▶ CPU

- 5.5 GHz
- Enhanced Out-Of-Order

▶ Caches

- L1 private 64k i, 96k d
- L2 private 1 MB i + 1 MB d
- L3 shared 48 MB / chip
- L4 shared 384 MB / book



Formulas – zEC12 / zBC12

Workload Characterization
L1 Sourcing from cache/memory hierarchy

Metric	Calculation – <i>note all fields are deltas between intervals</i>
CPI	$B0 / B1$
PRBSTATE	$(P33 / B1) * 100$
L1MP	$((B2+B4) / B1) * 100$
L2P	$((E130+E131+E132) / (B2+B4)) * 100$
L3P	$((E144+E150+E153+E159) / (B2+B4)) * 100$
L4LP	$((E147+E145+E151+E156+E154+E160) / (B2+B4)) * 100$
L4RP	$((E148+E146+E152+E157+E155+E161) / (B2+B4)) * 100$
MEMP	$((E135+E137) + (B2+B4-E130-E131-E132-E144-E150-E153-E159-E147-E145-E151-E156-E154-E160-E148-E146-E152-E157-E155-E161-E135-E137)) / (B2+B4)) * 100$
LPARCPU	$(((1/CPSP/1,000,000) * B0) / \text{Interval in Seconds}) * 100$

CPI – Cycles per Instruction

Prb State - % Problem State

L1MP – Level 1 Miss Per 100 instructions

L2P – % sourced from Level 2 cache

L3P – % sourced from Level 3 on same Chip cache

L4LP – % sourced from Level 4 Local cache (on same book)

L4RP – % sourced from Level 4 Remote cache (on different book)

MEMP - % sourced from Memory

LPARCPU - APPL% (GCPs, zAAPs, zIIPs) captured and uncaptured

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See “The Load-Program-Parameter and CPU-Measurement Facilities” SA23-2260-03 for full description

E* - Extended Counters - Counter Number

See “IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196, z114 and zEC12” SA23-2261-02 for full description

CPSP - SMF113_2_CPSP “CPU Speed”

Formulas – zEC12 / zBC12 Additional

Metric	Calculation – <i>note all fields are deltas between intervals</i>
Est Instr Cmplx CPI	CPI – Estimated Finite CPI
Est Finite CPI	$((B3+B5) / B1) * (.54 + (0.04*RNI))$
Est SCPL1M	$((B3+B5) / (B2+B4)) * (.54 + (0.04*RNI))$
Rel Nest Intensity	$2.3*(0.4*L3P + 1.2*L4LP + 2.7*L4RP + 8.2*MEMP) / 100$
Eff GHz	CPSP / 1000

Updated January 2015

Note these Formulas may change in the future

Est Instr Cmplx CPI – Estimated Instruction Complexity CPI (infinite L1)

Est Finite CPI – Estimated CPI from Finite cache/memory

Est SCPL1M – Estimated Sourcing Cycles per Level 1 Miss

Rel Nest Intensity – Reflects distribution and latency of sourcing from shared caches and memory


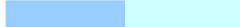
Eff GHz – Effective gigahertz for GCPs, cycles per nanosecond

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See “The Load-Program-Parameter and CPU-Measurement Facilities” SA23-2260-03 for full description

CPSP - SMF113_2_CPSP “CPU Speed”

 Workload Characterization
 L1 Sourcing from cache/memory hierarchy

Formulas zEC12 / zBC12 – Additional TLB

Metric – zEC12	Calculation – <i>note all fields are deltas between intervals</i>
Est. TLB1 CPU Miss % of Total CPU	$((E128+E129) / B0) * 100 * .65$
Estimated TLB1 Cycles per TLB Miss	$(E128+E129) / (E133+E140) * .65$
PTE % of all TLB1 Misses	$(E141 / (E133+E140)) * 100$

Note these Formulas may change in the future

Est. TLB1 CPU Miss % of Total CPU - Estimated TLB CPU % of Total CPU
 Estimated TLB1 Cycles per TLB Miss – Estimated Cycles per TLB Miss
 PTE % of all TLB1 Misses – Page Table Entry % misses

B* - Basic Counter Set - Counter Number

See “The Load-Program-Parameter and CPU-Measurement Facilities” SA23-2260-03 for full description

E* - Extended Counters - Counter Number

See “IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196, z114 and zEC12” SA23-2261-02 for full description

z10 and z196 Metrics

z196 versus z10 hardware comparison


■ z10 EC

▶ CPU

– 4.4 GHz

▶ Caches

- L1 private 64k i, 128k d
- L1.5 private 3 MB
- L2 shared 48 MB / book
- book interconnect: star



■ z196


▶ CPU

– 5.2 GHz


– Out-Of-Order execution

▶ Caches

- L1 private 64k i, 128k d
- L2 private 1.5 MB
- L3 shared 24 MB / chip
- L4 shared 192 MB / book
- book interconnect: star



Formulas – z10



Metric	Calculation – <i>note all fields are deltas between intervals</i>
CPI	$B0 / B1$
PRBSTATE	$(P33 / B1) * 100$
L1MP	$((B2+B4) / B1) * 100$
L15P	$((E128+E129) / (B2+B4)) * 100$
L2LP	$((E130+E131) / (B2+B4)) * 100$
L2RP	$((E132+E133) / (B2+B4)) * 100$
MEMP	$((E134+E135) + (B2+B4-E128-E129-E130-E131-E132-E133-E134-E135)) / (B2+B4) * 100$
LPARCPU	$(((1/CPSP/1,000,000) * B0) / \text{Interval in Seconds}) * 100$

CPI – Cycles per Instruction

PRBSTATE - % Problem State

L1MP – Level 1 Miss Per 100 instructions

L15P – % sourced from L1.5 cache

L2LP – % sourced from Level 2 Local cache (on same book)

L2RP – % sourced from Level 2 Remote cache (on different book)

MEMP - % sourced from Memory

LPARCPU - APPL% (GCPs, zAAPs, zIIPs) captured and uncaptured

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See “The Load-Program-Parameter and CPU-Measurement Facilities” SA23-2260-03 for full description

E* - Extended Counters - Counter Number

See “IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196, z114 and zEC12” SA23-2261-02 for full description

CPSP - SMF113_2_CPSP “CPU Speed”

Formulas – z10 Additional

Metric	Calculation – note all fields are <i>deltas</i> between intervals
Est Instr Cmplx CPI	CPI – Estimated Finite CPI
Est Finite CPI	$((B3+B5) / B1) * .84$
Est SCPL1M	$((B3+B5) / (B2+B4)) * .84$
Rel Nest Intensity	$(1.0*L2LP + 2.4*L2RP + 7.5*MEMP) / 100$
Eff GHz	CPSP / 1000

Note these Formulas may change in the future

Est Instr Cmplx CPI – Estimated Instruction Complexity CPI (infinite L1)

Est Finite CPI – Estimated CPI from Finite cache/memory

Est SCPL1M – Estimated Sourcing Cycles per Level 1 Miss

Rel Nest Intensity – Reflects distribution and latency of sourcing from shared caches and memory

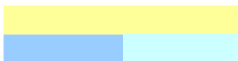
Eff GHz – Effective gigahertz for GCPs, cycles per nanosecond

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See "The Load-Program-Parameter and CPU-Measurement Facilities"
SA23-2260-03 for full description

CPSP - SMF113_2_CPSP "CPU Speed"

 Workload Characterization
L1 Sourcing from cache/memory hierarchy

Formulas – z196

Workload Characterization
L1 Sourcing from cache/memory hierarchy

Metric	Calculation – <i>note all fields are deltas between intervals</i>
CPI	$B0 / B1$
PRBSTATE	$(P33 / B1) * 100$
L1MP	$((B2+B4) / B1) * 100$
L2P	$((E128+E129) / (B2+B4)) * 100$
L3P	$((E150+E153) / (B2+B4)) * 100$
L4LP	$((E135+E136+E152+E155) / (B2+B4)) * 100$
L4RP	$((E138+E139+E134+E143) / (B2+B4)) * 100$
MEMP	$((E141+E142) + (B2+B4-E128-E129-E150-E153-E135-E136-E152-E155-E138-E139-E134-E143-E141-E142)) / (B2+B4)) * 100$
LPARCPU	$(((1/CPSP/1,000,000) * B0) / \text{Interval in Seconds}) * 100$

CPI – Cycles per Instruction

Prb State - % Problem State

L1MP – Level 1 Miss Per 100 instructions

L2P – % sourced from Level 2 cache

L3P – % sourced from Level 3 on same Chip cache

L4LP – % sourced from Level 4 Local cache (on same book)

L4RP – % sourced from Level 4 Remote cache (on different book)

MEMP - % sourced from Memory

LPARCPU - APPL% (GCPs, zAAPs, zIIPs) captured and uncaptured

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See “The Load-Program-Parameter and CPU-Measurement Facilities” SA23-2260-03 for full description

E* - Extended Counters - Counter Number

See “IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196, z114 and zEC12” SA23-2261-02 for full description

CPSP - SMF113_2_CPSP “CPU Speed”

Formulas – z196 Additional

Metric	Calculation – <i>note all fields are deltas between intervals</i>
Est Instr Cmplx CPI	CPI – Estimated Finite CPI
Est Finite CPI	$((B3+B5) / B1) * (.59 + (0.1*RNI))$ updated *
Est SCPL1M	$((B3+B5) / (B2+B4)) * (.59 + (0.1*RNI))$ updated *
Rel Nest Intensity	$1.67*(0.4*L3P + 1.0*L4LP + 2.4*L4RP + 7.5*MEMP) / 100$ updated *
Eff GHz	CPSP / 1000

Note these Formulas may change in the future
* Updated July 2012

Est Instr Cmplx CPI – Estimated Instruction Complexity CPI (infinite L1)

Est Finite CPI – Estimated CPI from Finite cache/memory

Est SCPL1M – Estimated Sourcing Cycles per Level 1 Miss

Rel Nest Intensity – Reflects distribution and latency of sourcing from shared caches and memory



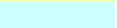
Eff GHz – Effective gigahertz for GCPs, cycles per nanosecond

B* - Basic Counter Set - Counter Number

P* - Problem-State Counter Set - Counter Number

See “The Load-Program-Parameter and CPU-Measurement Facilities”
SA23-2260-03 for full description

CPSP - SMF113_2_CPSP “CPU Speed”

 Workload Characterization
  L1 Sourcing from cache/memory hierarchy

Formulas – Additional TLB

Metric – z10	Calculation – note all fields are <i>deltas</i> between intervals
Est. TLB1 CPU Miss % of Total CPU	$((E145+E146) / B0) * 100 * .31 *$
Estimated TLB1 Cycles per TLB Miss	$(E145+E146) / (E138+E139) * .31 *$
PTE % of all TLB1 Misses	$(E140 / (E138+E139)) * 100$
Metric – z196	Calculation – note all fields are <i>deltas</i> between intervals
Est. TLB1 CPU Miss % of Total CPU	$((E130+E131) / B0) * 100 * .61 *$
Estimated TLB1 Cycles per TLB Miss	$(E130+E131) / (E144+E145) * .61 *$
PTE % of all TLB1 Misses	$(E146 / (E144+E145)) * 100$

Note these Formulas may change in the future

* Updated March 2012 / August 2012

Est. TLB1 CPU Miss % of Total CPU - Estimated TLB CPU % of Total CPU B* - Basic Counter Set - Counter Number

Estimated TLB1 Cycles per TLB Miss – Estimated Cycles per TLB Miss

PTE % of all TLB1 Misses – Page Table Entry % misses

See “The Load-Program-Parameter and CPU-Measurement Facilities” SA23-2260-03 for full description

E* - Extended Counters - Counter Number

See “IBM The CPU-Measurement Facility Extended Counters Definition for z10, z196, z114 and zEC12” SA23-2261-02 for full description

Definitions

CPI – Cycles per Instruction

PRB STATE - % Problem State

L1MP – Level 1 Miss Per 100 instructions

L15P / L2P – % sourced from L1.5 or L2 cache

L2LP – % sourced from Level 2 (or L4) Local cache (on same book)

L2RP – % sourced from Level 2 (or L4) Remote cache (on different book)

L3P – % sourced from L3 cache

MEMP - % sourced from Memory

LPARCPU - APPL% (GCPs, zAAPs, zIIPs) captured and uncaptured

Est Instr Cmplx CPI – Estimated Instruction Complexity CPI

Est Finite CPI - Estimated Finite CPI

Est SCPL1M – Estimated Sourcing Cycles per L1 Miss Per 100 instructions

Rel Nest Intensity – Relative Nest Intensity

Eff GHz – Effective Gigahertz

Machine Type – Machine Type (e.g. z10, z196, zEC12)

LSPR Wkld – LSPR Workload match based on L1MP and RNI

Pool – 1 = GCP, 3 = zAAP, 6 = zIIP

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by © are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a more complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*BladeCenter®, CICS®, DataPower®, DB2®, e business(logo)®, ESCON, eServer, FICON®, IBM®, IBM (logo)®, IMS, MVS, OS/390®, POWER6®, POWER6+, POWER7®, Power Architecture®, PowerVM®, PureFlex, PureSystems, S/390®, ServerProven®, Sysplex Timer®, System p®, System p5, System x®, z Systems®, System z9®, System z10®, WebSphere®, X-Architecture®, z13™, z Systems™, z9®, z10, z/Architecture®, z/OS®, z/VM®, z/VSE®, zEnterprise®, zSeries®

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.

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.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

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.

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

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

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.

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

* All other products may be trademarks or registered trademarks of their respective 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 Sync 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 Sync 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.

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SEs only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at:

www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT").

No other workload processing is authorized for execution on an SE.

IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.