Machine Type 2817 (A new mainframe generation)
- **5 Models** M15, M32, M49, M66 and M80
- **Processor Units (PUs)** 20 (24 for M80) PU cores/book
- Up to 14 SAPs per system, standard
- Dependant on the H/W model: Up to 15,32,49,66 or 80 PU cores available for characterization
- Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z Application Assist Processors (zAAPs), System z Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs) for I/O Operations.
- 2 spares designated per system
- Sub-capacity available for up to 15 CPs
  - 3 sub-capacity points (See lower left)

Memory System minimum of 32 GB
- Up to 768 GB per book
- Up to 3 TB for System and up to 1 TB per LPAR
  - 32/64/96/112/128 GB increments
- 16 GB Fixed HSA, standard
- I/O Up to 48 I/O Interconnects per System @ 6 GBps each
- Up to 4 Logical Channel Subsystems (LCSSs) and **STP**
- zNext is System z's 1st CMOS Out Of Order (**OOO**) core
- OOO yields significant performance benefit for compute intensive apps through *re-ordering instruction* execution (see illustration below)
  - Later (younger) instructions can execute ahead of a stalled instruction
  - *Re-ordering storage* accesses and parallel storage accesses
  - OOO maintains good performance growth for "traditional" applications

Vision: To deliver the best of all worlds mainframe, UNIX, x86 function processors integrated in a single system for ultimate simplicity, service and cost effectiveness across multiple heterogeneous workloads. The IBM zEnterprise™ System is a *first-of-a-kind* workload-optimized multiplatform (or multi-architecture) computing environment that spans (and tightly integrates) mainframe and distributed technologies.

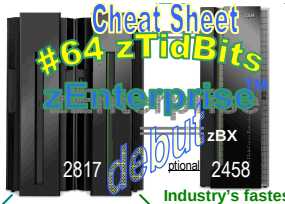

Up to Four active cores per chip
- 5.2 GHz
- L1 cache/ core
  - 64 KB I-cache
  - 128 KB D-cache
- 1.5 MB private L2 cache/ core

Two Co-processors (COP) Crypto & compression accelerators
- Includes 16KB cache
- Shared by two cores

24MB eDRAM L3 Cache
- Shared by all four cores

Interface to SC chip / L4 cache
- 40+ GB/sec to each of 2 SCs

I/O Bus Controller (GX)
- Interface to Host Channel Adapter (HCA)

Memory Controller (MC)
- Interface to controller on memory DIMMs
- Supports RAIM design

Over 100 new instructions

zEnterprise (z196)

Industry's fastest and most scalable enterprise server

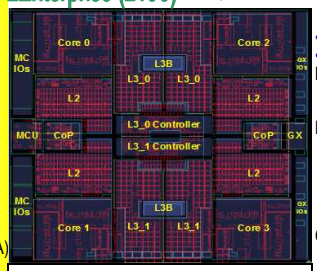

40% improvement for n-way processors
60% more total capacity than System z10

The integration of System z and distributed technologies into a revolutionary combination using the IBM Blade Center(zBX)

Integrated Hybrid System

- Runs applications unchanged and supports what you know using logical device integration between System z and distributed resources
- Workload specific accelerators to deliver significant performance and/or **lower cost per transaction**
- An ensemble is a collection of up to 8 z196 nodes that are managed collectively by the **Unified Resource Manager** as a single logical virtualized system
- A z196 **node** is a CPC with 0 or 1 zBX.
- The zBX may contain from 1 to 4 racks each containing up to 2 blade centers.
  NOTE: Nodes can all be CPCs with no zBXs.
- z196 nodes are deployed within a single site.
- Blade-based **fit-for-purpose** solutions
  POWER7 and x86 (SOD) IBM Server Blades will be supported
- Unified Resource Manager (zHMC)
- Virtual Resource Management and Automation

Machine Type/Model 2458-002
- 1 Model with 5 pre-configured Solutions for IBM Smart Analytics Optimizer (ISAOpt)

Racks – Up to 4 (B, C, D and E)
- 42U Enterprise, (36u height reduction option)
- 4 maximum, 2 chassis/rack 2-4 power line cords/rack
- Non-acoustic doors as standard
- Optional Acoustic Doors
- Optional Rear Door Heat Exchanger (conditioned water required)

Chassis – Up to 2 per rack
- 9U BladeCenter
- Redundant Power, cooling and management modules
- Network Modules
- I/O Modules

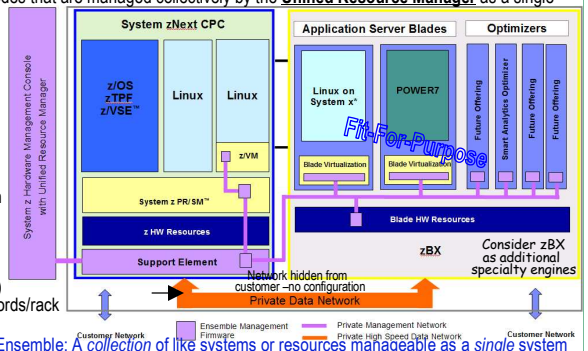

Ensemble: A *collection* of like systems or resources manageable as a *single* system

All Blades will be registered and certified for use by the zManager

Runs apps unchanged and supports what you know

Accelerator Blades:
- Smart Analytics Optimizer (SAO)
- Data Power Blades – provide deep-content routing and data aggregation (SOD)§
- Application-layer security and threat protection
- Protocol and message bridging
- Internet Security Services (iSS) appliances
- High-Performance Computing (HPC)

Blades (Maximum 112 in 4 racks)
- IBM Smart Analytic Optimizer Blades (up to 7 to 56)
- POWER7 Application Server Blades (up to 0 to 112)
- x Application Server Blades (up to 0 to 112)§

Management Firmware
- SE/HMC Hardware management

Fully configured 8 system cluster ~ 900 blades

Top of Rack (TOR) Switches - 4
- 1 GbE intra node management network (INMN)
- 10 GBe intra ensemble data network (IEDN)

Network and I/O Modules
- 1 GbE and 10 GbE modules
- 8 Gb Fibre Channel (FC) connected to customer supplied disks
- IBM Smart Analytic Optimizer uses DS5020 disks

DS5020 — Fibre Channel Disk Storage (customer responsibility)
Up to 112 drives in EXP520
8 Gbps FC or 1 Gbps iSCSI
2 or 4 GB of cache
Note: ISAO disk is not a SAN

Ensemble-related functions: zEnterprise™ Unified Resource Manager (a.k.a. zManager)
- The HMC will now be authoritative holder of some ensemble-scoped configuration not held by any of the Nodes in the ensemble
- Some configuration actions will be available ONLY from the HMC managing the ensemble, not the SE
- HMC will have a role in monitoring of workload performance
- Any V2.11.0 HMC can become either the Primary HMC (that controls the ensemble) or an Alternate (back up) HMC
- The Primary HMC can perform all non-ensemble HMC functions on CPCs that aren't members of the ensemble
- The HMC that performed the "Create Ensemble" wizard becomes the *only* Primary HMC
- The Primary HMC can also perform all non-ensemble HMC functions, even on CPCs that aren't members of the ensemble
- The Alternate HMC is specified when executing the "Create Ensemble" wizard after running the "Manage Alternate Hardware Management Console" task

NOTE: A Primary HMC is the *only* HMC that can perform ensemble related management tasks (create virtual server, manage virtual networks, create workload ....)

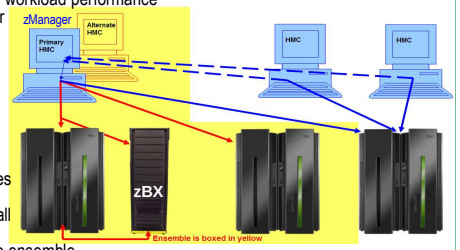

IBM Smart Analytics Optimizer is designed for OLAP-style SQL queries:
- Relational star schema (large **fact table** joined to multiple **dimensions**)
- Large subset of data warehouse accessed, reduced significantly by...
- **Aggregations** (SUM, AVG, COUNT) and optional **grouping** (GROUP BY)
- **Looking for trends or exceptions**

DB2 for z/OS routes SQL querires to accelerator transparently
- User need not change SQL or applications.

Appliance Optimizer up to 50x better performance

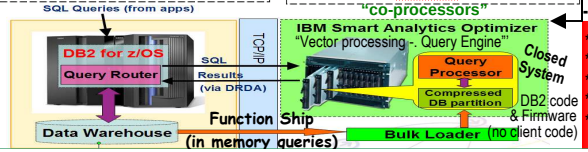

* Options to help in the elimination of hotspots and save energy
* Optional water cooling, high voltage DC power, top exit I/O
* Static Power Savings
* Query Max Potential Power
* Humidity and Altitude Sensors

HMC Set Power Saving

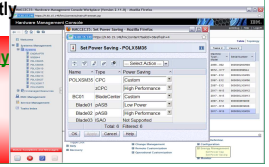

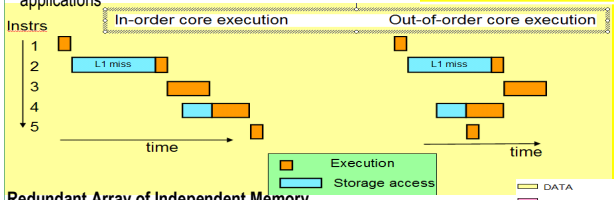

Redundant Array of Independent Memory
- 5 channel memory controller
- DIMM bus CRC error retry
- "Industry" leading reliability

Up to 3TB Memory capacity
- 3 MCUs per MCM
- 2-deep DIMM cascade

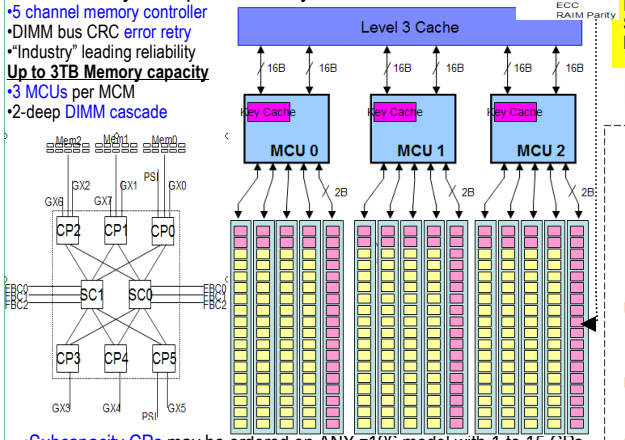

z10 MCM
96mm x 96mm in size
5 PU chips per MCM
Quad core chips: 3 or 4 active cores
PU Chip size 21.97 mm x 21.17 mm
4.4 GHz
Superscalar, In order execution
L1: 64K I /128K D private/core
L1.5: 3M I+D private/core
2 SC chips per MCM
L2: 2 x 24 M = 48 M L2 per book
SC Chip size 21.11 mm x 21.71 mm
Power 1800 Watts

comparison

zEnterprise (z196) MCM
96mm x 96mm in size
6 PU chips per MCM
Quad core chips:3 or 4 active cores
PU Chip size 23.5 mm x 21.8 mm
5.2 GHz
Superscalar, OOO execution
L1: 64K I / 128K D private/core
L2: 1.5M I+D private/core
L3: 24MB/chip – shared
2 SC chips per MCM
L4: 2 x 96 MB = 192 MB L4 per book
SC Chip size 24.4 mm x 219.6 mm
Power 1800 Watts

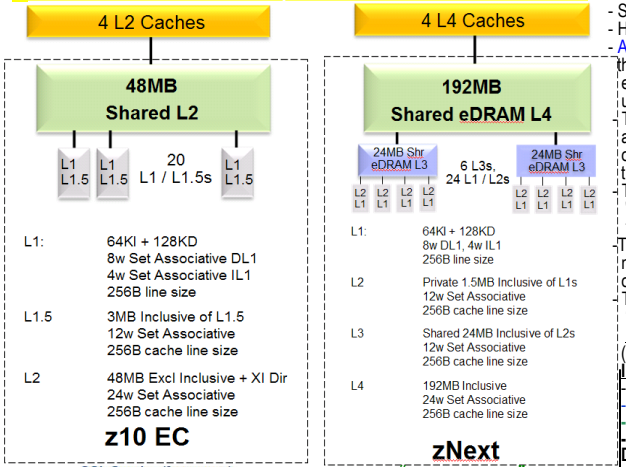

- Subcapacity CPs may be ordered on ANY z196 model with 1 to 15 CPs.
  - If 16 or more CPs are ordered all must be full 7xx capacity
  - All CPs on a zNext CPC must be the same capacity
  - All specialty engines run at full capacity. The one for one entitlement to purchase one zAAP and one zIIP for each CP purchased is the same for CPs of any capacity.
  - Only 15 CPs can have granular capacity but other PU cores may be characterized as full capacity specialty engines

CP MSU Capacity
Relative to Full Capacity
7xx = 100%
6xx ~ 64%
5xx ~ 49%
4xx ~ 20%
xx = 01 through 15

Three different sub-capacity levels, giving a total of 125 distinct capacity settings in the system, and providing for a range of over 1:200 in processing power.

§ Statement Of Direction

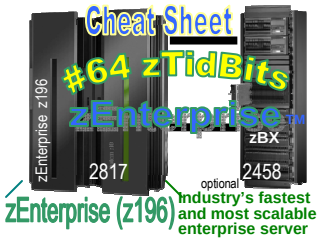
Cheat Sheet
#64 zTidBits
zEnterprise™
zEnterprise z196
zBX
2817
optional
2458
zEnterprise (z196)
Industry's fastest
and most scalable
enterprise server


zEnterprise z196

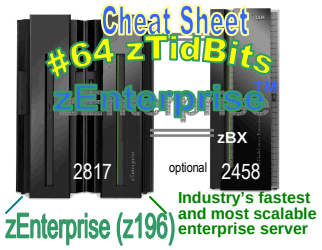
Cheat Sheet
#64 zTidBits
zEnterprise™
zBX
2817
optional
2458
zEnterprise (z196)
Industry's fastest
and most scalable
enterprise server