What's new…

How used….

When used…

What's new…	How used….	When used…
Did you know the z10 EC provides record level capacity over the previous System z servers that is achieved by both increasing the performance of the individual processor units and increasing the number of PUs per server. This is the first major redesign of instruction processing since 1997. The z10 EC book has a new Multi-chip Module (MCM) with five new IBM z10 Processor chips. Depending on the MCM version (17 PU or 20 PU), a z10 model can provide from 17 to 77 PUs, on one to four books characterizing up to 64 Pus by a customer. Z10 is estimated to provide up to 70% more total system capacity than the z9 EC Model S54.	The previous z9 outputted approximately 580 MIPS / CP where the z10 has a net increase of ~ 62% to 920 MIPS. This also pertains to any heritage purchased specialty processors at a no-charge upgrade. The jump in processor speed increased from 1.7GHz on z9 to 4.4GHz on z10. With this increase in processor cycles, the z10 is 85% the speed of System p (5.0GHz).	Products and their workloads will benefit in throughput and overall performance using this new quad chip design offering a greater opportunity for new and consolidated workloads to run collectively on a single System z platform. z10 EC introduces just-in-time deployment of capacity resources designed to provide more flexibility to dynamically change capacity when business requirements change. You are no longer limited by one offering configuration now having multiple configurations active at once and the configurations themselves are flexible so you can activate only what is needed from your defined configuration.
Did you know with the increase in CPU speed the new memory capacity has also increased from 512 GBs to 1520 GBs of memory. An increase of approximately 3 times the previous memory storage on the z9. In addition, we provide at no-fee fixed memory of 16 GB for your system configuration which in the past you paid for out of your storage purchase. The z10 provides free customer-memory to configure your systems.	We now have up to 384 GBs per book on the z10, previously on z9 we had a max of 128 GBs. The new z10 memory DIMMs plugged into 48 available slots, providing 64 GB to 384 GB of physical memory. The minimum memory in a book is 64 GB, installed in 16 DIMMs of 4 GB each. You can ensure application growth does not affect outage requirements due to system reconfiguration. The new no-fee system memory (16GBs) makes certain that there is less down time providing higher application availability.	With the increase in storage capacity, more applications can be consolidated in a single LPAR image for co- location to increase performance. Customers have more opportunity to place applications components in main memory above the bar – using z' 64 bit addressing. The locality of reference is minimized providing new design capabilities. New standard system memory now offers increased application up time, offering potentially more than the 99.999% in availability.
Did you know along with the new processor design the internal cache structure has a new improved access pattern and was rearchitected to offer greater speed accessing data and instructions. In fact, z10 now has 3 levels of cache (L1, L1.5, L2) rather than z9's two level cache (L1, L2) to maintain the new cycle speed. The z10s new L1.5 is a 3MB cache provided on each chip and includes a new a L2 48 MB upgrade from the z9s 40MB L2 cache.	The z9 used two concentric loops or rings (one flowing clock-wise ,the other flowing counter clock-wise) such as in a four-book system each book only is connected to two others, which meant that data transfer to the 3rd book required passing through one of the other books. The z10 uses Books are inter connected by a point-to-point (star) connection topology. This allows every book to communicate with every other book directly without intermediary.	The new cache topology enhances the overall performance for all applications and products executing on a z10. There is no latency to access cache data that resides on an alternate book from the processing instructions. This increases any workload's performance.
z/OS on the IBM System z10 implements a new approach to dispatching work to address client's increasing workload demand for processor cycles and high speed memory access. The new z/OS HiperDispatch makes improvements by having z/OS aware of the underlying physical topology of configured processors. z/OS can use this awareness to attempt to re-dispatch a unit of work repeatedly on the same physical CPU, or collection of physically adjacent CPUs (affinity node) to increase the chances of obtaining data from cache associated with those processors, instead of incurring a time delay by going outside of the local CPUs to an alternate book .	The mainframe used symmetric multi-processing (SMP) for many years where any unit of work can be dispatched on any logical processor. With the debut of specialty processors we instituted Asymmetric multi-processing where specialized processors are used for specific tasks such as zAAP and zIIP. Hiperdispatch now provides NUMA functionality where processors are grouped into nodes to create an affinity to their cache. This creates a tight working set running certain transactional type applications. Z10 uses 3 modes of dispatching: SMP, ASMP and NUMA.	This processing configuration flexibility can align application workload behavior to a customer's infrastructure design. This can enhance performance and overall throughput for certain applications. No other platform offers this form of agility to map workload characteristics based on a machine's hardware design. Very large z/OS images configured on groups of physical processors spanning multiple books and involving memory operations with intensive reference patterns are likely to achieve the improvements. I.E Transaction oriented application environments.
Did you know the z10 has a new architecture for its I/O bus. Past generation mainframes used a proprietary technology called Self-Timed Innerconnect (STI). The new z10 implementation uses the industry wide supported *Infiniband* and has allowed our I/O capability to more than double what it was on the z9. This new technology is designed to reduce overhead and latency, and provides increased data throughput.	InfiniBand is the result of merging two competing designs, Future I/O, developed by Compaq, IBM, and Hewlett-Packard, with Next Generation I/O developed by Intel, Microsoft, and Sun. From the Compaq side, the roots were derived from Tandem's ServerNet. InfiniBand is a switched fabric communications link used in high-performance computing. Its features include quality of service and failover, and it is designed to be scalable.	Infiniband establishes a new generation in I/O capability and throughput. In the previous system z9 an STI has data rates of 2.7 GBps. Infiniband technology uses a 6GBps data rate. This enhances overall performance of applications and products performing I/O such as DB2, IMS, VSAM, HFS, and zFS. Infiniband is also incorporated into the Parallel Sysplex Coupling Facility (PSIFB). This enhances the connectivity performance between z10 machines by doubling the speed.

Did you know many customer data centers are now 10 to 15 years old, and the cooling facilities are not adapted to the present needs. Traditional cooling methods allowed for 2-3 kW of cooling per server rack. Today's requirements are 20-30 kW per rack. Heat density is many times past the design point of the data center. Did you know IBM is continually designing the mainframe for utility efficiency where the System z10 is 15% more energy efficient than z9 EC. The mainframe has had 16x energy improvement over last 12 years and still accelerating. IBM has improved the performance per watt for 10 CMOS generations. The z10 EC uses approximately 0.91 watts per MIPS where in the past the BiPolar system used approximately 350 watts per MIPS. Today's z10 has its own power and thermal information displayed via the System Activity Display (SAD) which is part of its hardware configuration console allowing you to monitor the mainframe's power usage.

The power and cooling issues data center managers face are not standalone challenges, they have a cascading impact on other facilities issues such as wiring, floor space, etc. If energy and cooling are not adequate, it can lead to a requirement for new facilities or major renovations of an existing data centers. For instance the cost to build a new data center with top of the line power capabilities can run over $1000 per square foot. Refurbishing costs of existing data centers can also prove prohibitive, such as new cooling units requiring reconfigurations of floors.

Did you know that distributed servers that their CPU utilization is maxed out between 15-18%, but that their energy consumption is still that of a fully loaded server.

IBM internally used the mainframe to achieve significantly utility savings using less hardware and consolidated 3,900 distributed servers going to 30 System z EC IFL servers resulting in +80% energy reduction, +85% space reduction.

Distributed platforms results in power and cooling requirements per square foot that stress current power thresholds. Because non-mainframe platforms have an Initial attractive price point their popularity has Increased, but the same time their heat has created a problem for data centers whose total power limit is consumed entirely by these servers. Distributed business users often select platforms on the basis of a singular application need, where tendency is to buy more hardware as these servers proliferate.

The costs of power over time also needs to be considered as part of data center planning. Along with the rising trends in energy costs there is still a continuing trend toward distributed servers which stress the power capacity of today's utility facilities. The z10 is positioned as an energy efficient platform that can save customers in overall utility expenses by initially 15%.

Did you know from a cost perspective, in place of z9s, that 33 z9 mainframes might be reduced to approximately 18 z10 EC mainframes, for even greater savings in IT running costs.

Did you know that 17 z10 EC mainframes with 64 IFLs (1,088 IFLs) each would replace the 23,000 distributed processor cores at approximately a 20 to 1 ratio.

Misperceptions arising from inaccurate charge-back methodologies can provoke a naïve, pre judgment that System z is more expensive than distributed servers. The mainframe IS more expensive than ONE distributed server, but not MANY distributed servers can establish what the exact numerical equivalency is for a single mainframe.

Sometimes customers and vendors fail to configure sufficient distributed capacity to meet the level of the mainframe. However, when studies are performed by consultants to establish true equivalence, critical evaluation will lead to economic advantages in favor of System z

When proper equivalence is established in terms of capacity and capability, System z10's Total Cost of Ownership will typically be lower than other platforms. But if comparisons are made without understanding the fundamental differences between the mainframe and distributed platforms, it is likely that false equivalences will be made that in turn yield incorrect results. The key differentiator of the mainframe – its capability of running high volumes of mixed transactions concurrently at high utilizations with guaranteed service levels, sharing every component (hardware software, labor, power, space…) – allows enterprises to exploit System z10 to run their essential value-generating computing at the lowest cost.

jkettner@us.ibm.com