

i-BigBlue

PROFESSIONALS' MONTHLY

I-BUSINESS AND INTERNETWORKING INVOLVING MAINFRAMES AND AS/400S

May 2003

HEADLINES

IBM's T-Rex, the codename for the mighty 32-way z990 mainframe, the 'meanest' mainframe to-date, shows how patently untrue the "dinosaur myth" was.

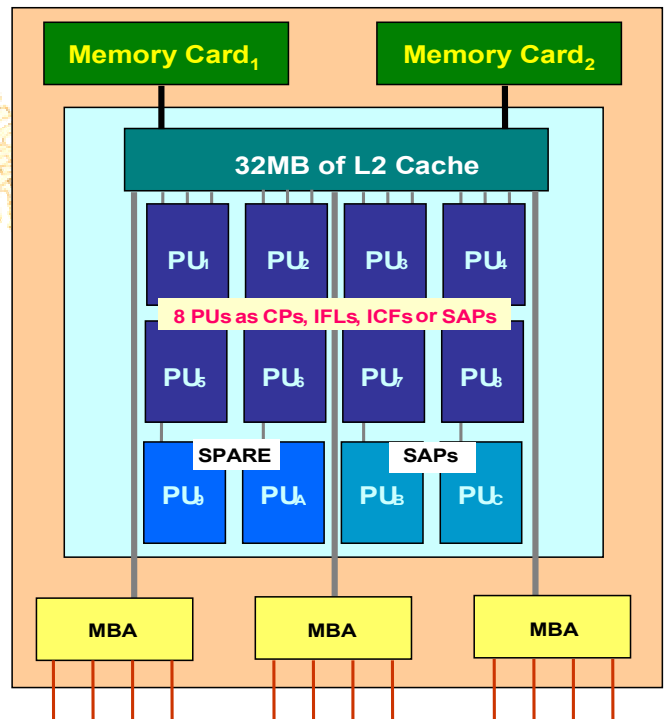

The z990 sets new thresholds for many key mainframe parameters; viz. 30 LPARs (from the hitherto 15), more than 16 CPs, 512 channels and 16 HiperSockets.

IBM yet again breaks its own record for SSL handshakes per second with the new z990 which can do 11,000 compared to the 7,000 record set last September with z/OS 1.4 running on a z900.

The 1.3GHz POWER4 was fast – the 1.45GHz POWER4+ that was unveiled last November was even faster. Now there is a 1.7GHz POWER4+ and you can have 32 of those in a p690 – along with 512GB of memory if you need a very heavy-duty Unix server.

While XML adoption still lags way behind prior expectations we have some interesting dialects, in particular MathML, just in case you wanted to represent mathematical notations using XML rather than the old fashioned way.

While IBM pushes the performance envelope at the high end of the server market with the z990, p690 and p670, Sun turns its attention to Intel powered 32-bit servers.



A z990 'book' – IBM's new packaging unit for processors that contain an MCM (in this instance a 12 processor MCM), cache memory, main memory and memory bus adapters for I/O.

CONTENTS

IBM's new z990 mainframe	2
POWER4+ driven enhancements to the IBM pSerers	11
XML 'exotica' – MathML?	13
Knowledge management with corporate portals: Final Part	14
i-sight: i-nteresting i-business i-nfo: by Deanna Gurugé	16

PUBLISHED BY: ANURA & DEANNA GURUGÉ


The Mighty, On-Demand' @server zSeries 990

If you stop and think about it, you would have to admit that it takes a special type of 'brass-plated' chutzpah to codename a multi-billion dollar flagship product after a dinosaur – even if it is indeed the king of dinosaurs, the rapaciously marauding, all powerful *Tyrannosaurus rex* (T-Rex). But in keeping with its stirring theme of naming new mainframes after conspicuous predators [e.g. raptor for the entry-level z800, Feb. 2002], IBM did just that with the all new z990 mainframe that it introduced on May 13.

However, unlike with the eLiza name (and reptilian logo) [April 2003], there is method in IBM's apparent folly. This name is a jab, à la "*The Lost World: Jurassic Park*", at IBM's server-side competition [which these days is limited to Sun and H-P] who will have no choice but to stay well clear of this new beast.

This new mainframe family, which again stretches the performance and capacity envelope for this genre, is the ultimate vindication of the supposed 'dinosaur myth'. Rather than facing extinction at the hands of 'low-cost' RISC Unix workstations, mainframes are still in business, and in demand – and now, moreover, poaching strategic business from Unix machines with their unparalleled Linux 'virtualization' capabilities.

As one who was in the front-lines at the height of the '*mainframes are dinosaurs*' war of the late 1980s, on the then beleaguered and oft ridiculed mainframe side, it is indeed invigorating to see the new frontiers being forged with the z990 – particularly given that exactly week prior to the z990 unveiling [i.e. May 6] IBM set yet another milestone in the Unix server side with the 1.7GHz POWER4+-based p690. Thanks to Lou, IBM is teaching dinosaurs how to dance.



The new IBM @server zSeries 990 in the two frame configuration – the A-Frame that holds the processors, memory and control hardware [i.e. the so called Central Electronic Complex (CEC)] and one I/O cage on the right, with the optional Z-Frame that can house two more I/O cages on the left.

source: IBM

The real noteworthy technical breakthroughs vis-à-vis mainframe computing available with the z990, which represents the first true major enhancement to the 64-bit z900 family which was introduced on October 3, 2000 [Oct. 2000] (not counting the 'Gen 1.5' performance upgrade in July 2002), include:

1. up to **32 processors** [i.e. 32-way] – double that of the 16-way available with the z900 [and up 20 from the max. 12-way machines available as recently as the year 2000 with S/390 G6 machines].
2. up to **30 Logical Partitions (LPARs)** – double the 15 available up until now.
3. **512 channels** per system, via two 256 channel *Logical Channel SubSystems (LCSSs)*, thus doubling the 256 channel limit of prior systems.


4. up to **256GB** of **memory** – 4x the 64GB max. on the z900.

5. **96GBytes** of **I/O bandwidth**, via up to 48 2GB/sec. Self-Timed Interconnect (STI) buses [3/26/1999], 4x the 24GB/sec. available on the z900. It is sobering to think that when we first encountered the STI in 1999 its max. speed was 333MB/sec – 1/6th of today's speed.

6. Ability to **purchase individual processors** from one to 32.

7. **48 OSA-Express ports** [as opposed to the 24 on z900s] giving 48Gbps of Ethernet network connectivity per system with the Gigabit Ethernet adapters.

8. **16 HiperSockets** IP-based 'internal



A graph from IBM showing the degree to which the z990 pushes out the envelope, compared to the z900, S/390 G6 and S/390 G5, when it comes to the number of CPUs, max. memory, processor cycle time and system I/O bandwidth.

source: IBM

memory-to-memory LANs inside your mainframe' [Sept. 2001 & diagram on page 5] for inter-LPAR communication. The previously limit was 4. It is worth noting that HiperSockets, a highly strategic technology for partitioned

	Intro.	MIPS per CPU	Δ	Max. No. CPUs per machine	Max. MIPS/ Machine	Max. Memory	Max. LPAR	Max. OSA-Express	Max. 800Mbps FICON
G1	Sept. '94	11 to 13		6-way	60		10	0	
G2	June '95	22	83%	10-way	165		10	0	
G3	Sept. '96	45	105%	10-way	325		10	0	
G4	June '97	63	40%	10-way	450		15	0	
G5	Aug. '98	152	141%	10-way	1,069	24GB	15	12	12
G6	May '99	201	32%	12-way	1,614	32GB	15	12	24
900	Oct. 2000	225(?)		16-way	2,500	64GB	15	24	96
900 Gen 1.5	April 2002	270(?)	20%	16-way	2,925	64GB	15	24	96
990	May 2003	410(?)	52%	32-way	8,134	256GB	30	48	120
800	Feb. 2002	185(?)		4-way	625	32GB	15	24	32

machines, has only been available since October 2001 with the **z/OS Ver. 1 Rel. 2**.

9. **On/Off Capacity on Demand** (On/Off CoD) that permits individual CPUs to be activated on a 24 hour billing-period basis.

10. **CP Assist for Cryptographic Function (CPACF)** on every processor – which among other things have allowed IBM to yet again set another record for **SSL handshakes** [Aug. 2002] by breaking its own record of 7,000 set with z/OS Ver. 1 Rel. 4 with 11,000 handshakes/second on a z990 – a 57% increase.

11. 64-bit integer instructions with 64-bit operands.

Only 4 Models

Though spanning more processors than ever before [i.e. 1-way to 32-way], IBM exploiting the possibilities opened up by *Capacity on Demand* features (such as the new On/Off CoD) is only offering four (4) models of the z990; viz. A08, B16, C24 and D32. In marked contrast there are 25 z900 models ranging from the 101 to 2C9 as was shown in July 2002. In the case of the z990, the two digits in the model number indicate max. CPUs possible on that model.



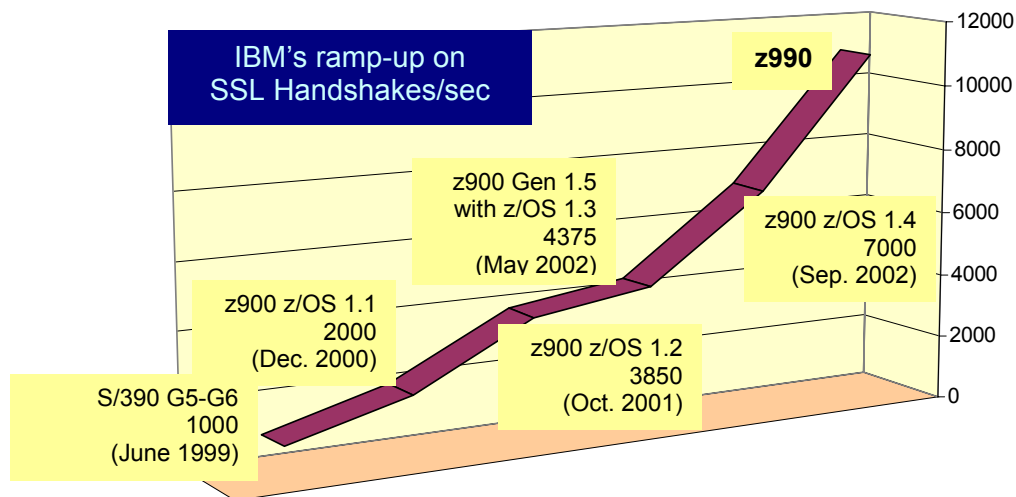
The z900 with its 'asymmetrical' frames on the left – and the new z990 with its 'balanced' frames on the right.

source: IBM



IBM's twin-towers of power, the p690 on the left and the z990 on the right (making one rather nostalgic for the ol' days when IBM had a RPQ for **Ffyes banana yellow** for mainframes given that Ffyes always ordered their mainframes in their trademark, bright yellow).

source: IBM



THE FOUR z990 MODELS							
z990 Model	Processor Books	'Business' Processor Units	System Assist Processors (SAPs)	Spare Processors	Memory	STI Buses	Max. I/O Cages
A08	I	8	2	2	8 – 64GB	12	Three (3)
B16	II	16	4	4	8 – 128GB	24	
C24	III	24	6	6	8 – 192GB	36	
D32	IV	32	8	8	8 – 256GB	48	

Doing it by the Book


The 'book', a notion initially introduced with the POWER4+-based p630 earlier this year [Feb. 2003], has now become IBM's new high-level processor packaging unit for post-2002 machines. A book is an integrated unit that contains one or more processors, associated memory and the necessary I/O bus adapters – where the memory in this case includes both cache and main memory. It is essentially a self-contained processor complex. The MultiChip Module (MCM) [or the SingleChip Module (SCM) in the case of smaller systems such as the p360], hitherto the basis of modularity for server systems, now becomes a component within a book.

The z990 is totally book-based. It is the number of books per system that differentiates the four z990 models from each other – and dictates the maximum number of processors, memory and I/O bus connections you can have on each model. In essence, each higher level model holds one additional book – as shown in the table above. However, since a book can contain from 8GB to 64GB of memory, increments of 8GB, the exact memory configuration of each model can still vary.

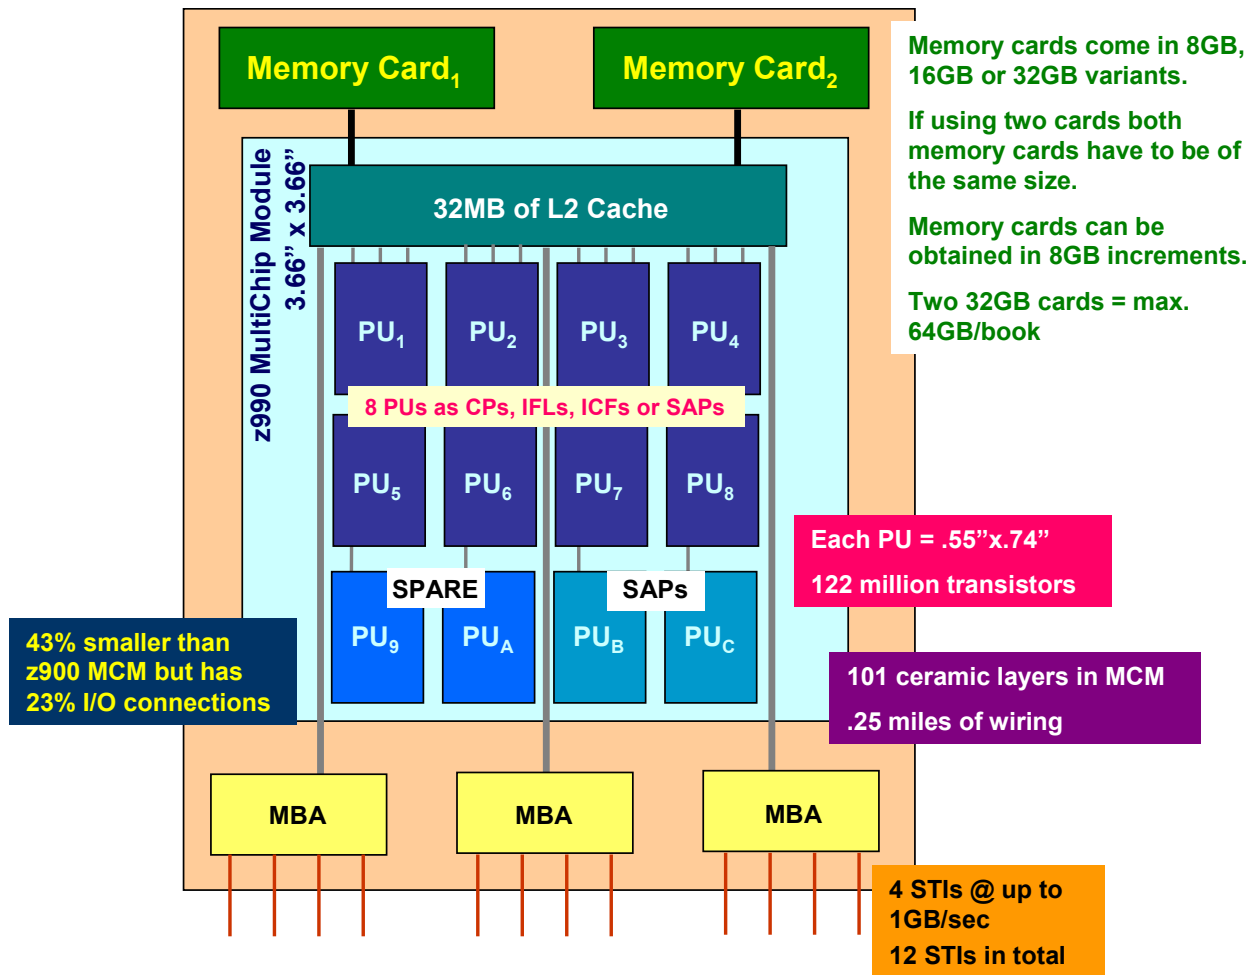
Each z990 book consists of:

1. one (1) 12 processor MCM with 32MB of built-in level 2 cache,

2. two (2) memory card adapters, and
3. three (3) Memory Bus Adapters (MBAs) – each of which supports four (4) Self-Timed Interconnect (STIs) buses thus providing a total of 12 STIs per book.



Another original IBM picture showing how HiperSockets relates to standard Sysplex inter-system communications via Cross-System Coupling Facility links [e.g. 800Mbps fiber optic HiPerLinks that can extend up to 24 miles.]



The high-level layout of the ultra-dense z990 book, which contains one 12 PU MCM, 2 memory cards and 3 Memory Bus Adapters (MBAs) for I/O connections.

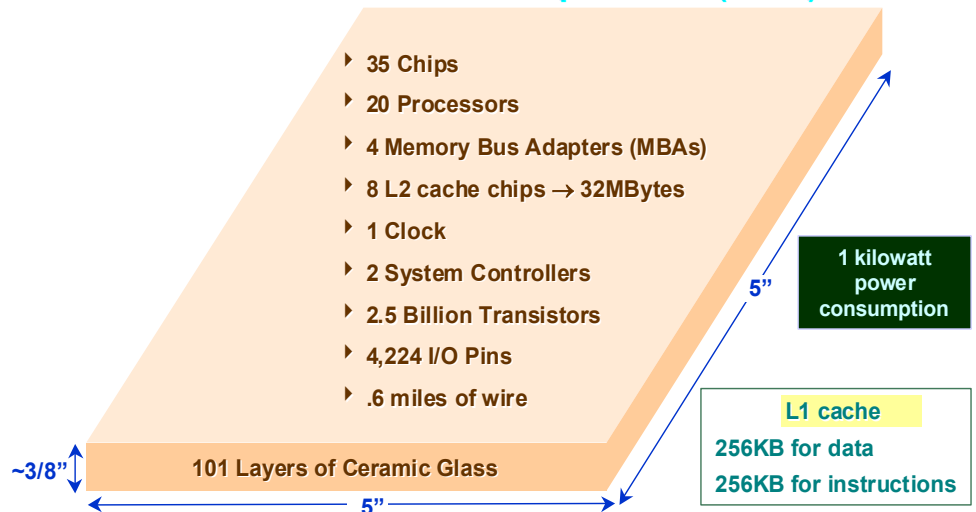
A z990 Model A08 has one book, the B16 has two, the C24 three, while the D32 has 4.

The memory cards for the z990 books come in 8GB, 16GB or 32GB variants. The memory cards can be purchased in 8GB increments. However, if two cards are to be installed in a single book, then they both have to be of the same size. The maximum 64GB of memory per book is achieved by using two 32GB memory cards.

The z990 MCM contains 12 PUs. This, interestingly, is down

from the 20 PUs found on a z900 MCM –

zSeries 900 Multi Chip Module (MCM)



though to be fair, the z990 MCM is 43% smaller (though still made up of 101 ceramic layers) and the processors are faster [i.e. 0.83 nanosecond cycle time versus 1.3 nanoseconds on the z900].

There has also been a fundamental change in the way the processors are deployed. In the z900s (and the G5/G6s for that matter) you could only have one MCM per machine. That immediately capped the number of processors per machine to being what was available on a single MCM -- less those reserved for backup purposes. Hence, the 16-way limit.

But now you can have multiple books per system -- and hence multiple MCMs per system (as is the case with the new i- and pSeries as discussed in [Jan. 2003](#) & [Oct. 2002](#)). Thus the new 32-way machines. In essence you can now have a full-blown parallel sysplex configuration within a single box -- particularly given that z990s only work in 'partitioned' LPAR mode.

Of the 12 PUs on each z990 MCM, 8 are available to be readily used as 'business'

processors, either as general purpose Central Processors (CPs), dedicated **Integrated Facility for Linux (IFL)** units, **Integrated Coupling Facility (ICF)** processors running the **Coupling Facility Control Code (CFCC)**. They, can if needed, also be configured to be used as additional **System Assist Processors (SAPs)**.

Of the other 4 PUs, two are always reserved for use as SAPs with the other being reserved as spares. Within this framework, it is possible to get z990 machines that have from 1 to 32 'business' processors -- in one processor increments. Thus, there are 32 model variants, spanning the four main models. These go from the single processor 301 to the 32-way 332.

However, with the z990 set-up you are not restricted in terms of what model you should order depending on how many PUs you plan to start off with. Thus, if you expect steady growth over a couple of years you could still start off with a up to 16-way Model B016 even if you initially only plan to use 5 business processors. This is where all of IBM's 'On Demand' offerings come into play.



Behind the covers of a 2-frame [i.e. A- and Z-Frame] z990 and a diagram of how the main A-Frame holds the CEC and the 1st I/O cage while the z-Frame is used to house I/O cages two and three.

source: IBM

z990 - OPERATING SYSTEM SUPPORT All in LPAR Mode

Compatibility Mode:

1. up to 15 LPARs
2. single Logical Channel Subsystem (LCSS); i.e. max. 256 channels

Operating System	Versions	Availability
OS/390	Ver. 2 Rel. 10	June 2003
z/OS	Ver 1 Rel 2, Rel 3 & Rel. 4	
z/VM	Ver 3.1.0 Ver 4.2.0, 4.3.0	
Linux for zSeries	Red Hat 7.1 & 7.2 SuSE SLES7 & SLES8	
VSE/ESA	2.5, 2.6, 2.7	
TPF/ESA	4.1	

Exploitation Mode:

1. up to 30 LPARs
2. two Logical Channel Subsystem (LCSS); i.e. max. 512 channels

Operating System	Versions	Availability
OS/390	N/A	
z/OS	Ver. 1 Rel. 4	October 2003
z/VM	4.4.0	August 2003
Linux for zSeries	Red Hat 7.1 & 7.2 SuSE SLES7, SLES8 & SLES9	4Q2003
VSE/ESA	2.5, 2.6, 2.7	October 2003
TPF/ESA	N/A	

The spare processor in a z990 book are, however, not meant to be used for any type of 'CUoD' activity. Instead, they are there purely to act as failover units in the event (however, remote) of an unexpected, catastrophic failure of any of the other PUs. This is worth remembering when looking at the capacity requirements of a new z990 since 'sparing' on a z900 was different in that you could get down to just 1 spare per system.

The new "On/Off CoD" feature provides an easy-to-enable capability to "turn on and turn off" CPs to meet fluctuating workload demands. A similar capability is available on z900s albeit only via an RPQ. The goal here is to allow customers to purchase capacity for future use without incurring the ongoing cost of increased software billings in the short-term. The CP activation and deactivations, moreover, can be totally non-disruptive.

STSI Instruction

In marked contrast to IBM mainframes prior to this, it is no longer possible with the z990 to determine the number of active PUs in a machine just from its model number -- i.e. A08, B16, C24, D32 -- since each model now offers a wide range of PU options. In order to facilitate software billing, there will now be a "software" model associated with the number of PUs that are characterized as CPs. This number will be obtained through the use of the *Store System Information (STSI)* instruction

There will be no affinity between the hardware model and the number of active CPs within a

Statements of Direction

1. Support for up to 60 LPARs in future.
2. up to four (4) Logical Channel Subsystems (LCSSs) providing a maximum of 1,024 channels.

machine. For example, it would be possible to have a model C24 which has 8 PUs characterized as CPs. For software billing purposes, the STSI instruction would report 308. Since this represents a significant 'over-booking', it would be more normal for a customer to get a B16 if all they want initially is just 8 active CPs. The STSI for such a

General Availability in Two [or Two-and-Half] Phases

1. z990 GA1 – availability as of June 16th, 2003

- ❖ Two low-end models only; viz. A08 and B16
- ❖ single Logical Channel Subsystem (LCSS); i.e. max. 256 channels
- ❖ up to 15 LPARs
- ❖ Compatibility mode only with all the Operating Systems (per page 8)
- ❖ Upgradeable from z900

2. z/VM V4 R4 – availability as of August 15th, 2003

3. z990 GA2 – availability as of October 31st, 2003

- ❖ Two high-end models; viz. C24 and D32
- ❖ Dual Logical Channel Subsystem (LCSS); i.e. max. 512 channels
- ❖ up to 30 LPARs
- ❖ new PCIX Cryptographic Coprocessor that replaces the current PCI and CMOS Cryptographic Coprocessors
- ❖ Exploitation mode with z/OS Ver. 1 Rel. 4
- ❖ concurrent model upgrades between the z990 models

machine would also say 308. This would also be the case if you were using a A08 or D32 with just 8 active CPs.

z990 On-Demand

Within reason, all of IBM's 'on-demand' features for the z990 work under the umbrella of IBM's unique and strategic **Customer Initiated Upgrade (CIU)** facility cum agreement/contract. CIU enables a customer to initiate a processor or memory upgrade provided that there are spare PUs available (not counting the two failover PUs included in each book) and there is already installed but unused memory in the machine.

CIU is initiated over the Web using IBM Resource Link. The upgrade, if feasible and approved by IBM, is downloaded and 'activated' using Remote Support Facility of the z990 Hardware Management Console.

The 'On/Off CoD', like all other 'on-demand' features falls under the purview of CIU. 'On/Off CoD' though useful is not as flexible as some of IBM's other options.

'On/Off CoD' is offered with z990 processors to deliver a temporary increase in capacity to address spikes in customer workload processing. It only allows a customer to temporarily turn on unassigned or unowned PUs available within a customer's current machine for use as CPs.

'On/Off CoD' does not support the temporary activation of IFLs, CFs, memory or channels – even if spare capacity was installed and is available. In order to use 'On/Off CoD' customers have to have a valid signed agreement for CIU.

Capacity Upgrade on Demand (CUoD), which has now been around for a few years on the mainframe front, is still a more powerful and flexible option than 'On/Off CoD' *per se*. CUoD allows for the nondisruptive addition of one or more CPs (à la 'On/Off CoD') as well as

ICFs or IFLs. CUoD will enable customers to quickly add processors up to the maximum number of available inactive PUs in their machine. It can also be used to dynamically upgrade all I/O cards in the I/O cages. CUoD combined with Parallel Sysplex technology enables mainframe enables virtually unlimited capacity upgrade capability.

Bottom Line

Whichever way you try to slice-and-dice this, it is yet another superlative *tour-de-force* by IBM's eServer hardware folks. I am not sure exactly what Lou did to them but they seem to be inspired and on intellectual steroids. From a raw performance standpoint the z990 reaffirms what mainframe computing is all about.

Rather than being dinosaurs, mainframe computing, thanks to IBM's concerted efforts, is healthier than it has been in quite a long time. IBM is claiming significant [i.e. greater than 50%] increases in new mainframe workloads. And that is what counts. According to IBM, mainframes gained 3% of worldwide market share in 2002 – though unfortunately these figures probably have their origins with the Incorrect Data Corp. (IDC). But these two numbers, straight from IBM, are gratifying: 75 new mainframe customers in 2001 and 100 in 2002.


The z990 sure pushes the envelope.

Given that, I am sure glad that IBM insists that z990s only work in LPAR mode. If not, I would have been concerned about the diminishing returns of SMP – something that IBM has drilled into all of us over the years. With z990 in LPAR mode what we are essentially getting is Parallel Sysplex within a box. Yes that was always possible, but now with the 16- to 32-way machines it makes even more sense. You can't fault these machines. With the continued economic slowdown, I am not sure we have the demand for the high-end as yet. But the good news is that it is there ... ready and waiting.

POWER4+ Driven Enhancements to the IBM pSeries

Since October 2001, IBM, emboldened by the redoubtable POWER4(+) [Oct 2002 to Feb. 2003] has been playing a taunting cat-and-mouse game with Sun. Every time Sun makes what it hopes is a significant move, IBM over-trumps it with an announcement that overshadows Sun in both performance and price.

System	SPECjbb 2000 Java instructions/sec
32-way , 1.3GHz IBM p690	339,484
72-way , 900Mhz Sun Sun Fire 15K with JVM 1.3.1 on Solaris	324,309
72-way, 900Mhz Sun Sun Fire 15K with JVM 1.4.0 on Solaris	433,166



Why the 32-way, p690 'Regatta' with 1.3GHz POWER4 processors eclipsed Sun's much vaunted 'Starcat' Sun Fire 15K last October when the two were announced 9 days apart.

SNAPSHOT OF IBM'S 'GENERAL PURPOSE' P SERIES POWER(+) PLAY

	p690	p670	p650	Enhanced p630
Processor Type	Old: 64-bit POWER4 New: POWER4+	Old: 64-bit POWER4 New: POWER4+	64-bit POWER4+	64-bit POWER4 or POWER4+
Max. # of processors	32	16	8	4
Max. clock speed	1.1 or 1.3GHz 1.5 & 1.7 - POWER4+	1.1GHz 1.5 - POWER4+	1.2 or 1.45GHz	1.0GHz with POWER4 1.45GHz POWER4+
n-way options	8, 16, 24 & 32	4, 8 & 16	2, 4, 6 & 8	1, 2 & 4
Data/Instruction Level 1 cache	32KB – 64KB	32KB – 64KB	32KB – 64KB	32KB – 64KB
Level 2 cache	5.7MB	5.7MB	1.5MB	1.44MB
Level 3 cache	128MB	128MB	8 or 32MB	32MB
Min/Max. memory	8GB/256GB new: 8GB/512GB	4GB/128GB	2GB/64GB	1GB/32GB
LPARs	16	16	8	4
Max. PCI slots	160	60	55	4 with POWER4 6 with POWER4+
PCI bus speed	66 or 33MHz	66 or 33MHz	133MHz	133Mhz with POWER4+ 33 with POWER4
Start prices	\$450,000	\$175,000	\$29,995	\$12,495

The 1.3GHz POWER4 as initially used in the p690 was plenty fast enough, as clearly demonstrated by the [SPECjbb 2000](#) chart on [page 10](#) – with a 32-way p690 being more brawny than a 72-way Sun Fire 15K. Then in November of last year we got the POWER4+, which was even faster.

The [SPECint 2000](#) ratings gives us a clue as to how fast the POWER4+ could be given that the SPECint for a 1.45GHz POWER4+ comes in at 909 versus 804 for the 1.3GHz POWER4. And that was at 1.45GHz.

The POWER4+ made its debut on the [p650](#) [[Nov. 2002](#)] – where the p650 is a mid-range machine that only goes up to 8-ways. Then in February of this year [[Feb. 2003](#)] IBM enhanced the p630 (and low-end machine that can only have up to 4 processors) with the POWER4+ -- again at 1.45GHz.

What we had been waiting was for IBM to unleash the POWER4+ on the 32-way, flagship p690. Well, IBM just did that on May 6th – just a week prior to unfurling the [z990](#) – just in case anybody still had any lingering doubts about its almighty clout in the high-end server business.

But IBM didn't put the 1.45GHz POWER4+ in the p690. Instead they opted for a even faster 1.7GHz version coupled with new 567MHz memory. IBM claims that a new 1.7GHz p690 is 65% faster than the original 1.3GHz p690. Given that the original p690 was way ahead of the Sun Fire 15K, this major boost in performance should put IBM so far ahead that Sun is unlikely to be even in range of shadow being cast by the p690.

In addition to the 1.7GHz version, IBM has also introduced a 'lower-cost' 1.5GHz POWER4+. This processor option is also now available with the p690 akin to the 1.1GHz and 1.3GHz options available on the original offering. IBM has also doubled the memory capacity of a p690 from 256GB to 512GB. As with the p630 and the new [z990](#), POWER4+ processors, cache memory and main memory are packaged within an integrated book.

The I/O bandwidth capabilities of the p690 have also been re-jigged with faster PCI-X planers. Overall, IBM claims that the peak I/O bandwidth capability of a p690 has been increased from 16GB/sec to 44GB/sec – a near three-fold jump.

IBM has also made some improvements to the CUoD capabilities of the p690 [[Jan. 2003](#)]. For a start, IBM's has made the available processor granularity that much better. With the old 'model' the minimum CUoD configuration possible was 12 processors – obviously spread across two MCMs, given that the maximum PUs you can get with a POWER4(+) MCM is 8. Now it is possible to have a minimum CUoD configuration with just 8 active PUs to begin with – even though you still need to order a minimum of two 8-way 1.5GHz or 1.7GHz MCMs.

As with the [z990s](#), the p690 also now has the "On/Off CoD" feature which enables pairs of p690 PUs to be activated, on a fee basis, for up to 60 'processor days' of usage – where a 'processor day' is measured each time a PU is activated for a 24-hour period or a part thereof.

p655 and p670 Enhancements

The 1.5GHz and 1.7GHz POWER4+ processors are now also available on the cluster-oriented, 'ultra-dense' p655 [i.e. 64 p655s grouped together in a pSeries 1600 Cluster] [[Nov. 2002](#)] – boosting the performance of these boxes as much as 83% per IBM. The p655 can also be partitioned – with up to 4 LPARs, with it being possible to run Linux on an LPAR.

The mid-range, 16-way capable p670 also gets the POWER4+ -- but just the 1.5GHz version. But this, nonetheless, is a big jump for the p670 given that it only had the 1.1GHz PU before. IBM obviously is trying to maintain a differentiation between the p690 and p670 by not offering the higher clock speed PUs on the p670.

Bottom line – WOW!

XML 'Exotica' – MathML?

To be fair, despite the ongoing media publicity (albeit now somewhat more subdued than it was a couple of years ago), it would appear that the jury is still out when it comes to XML. Though it is irrefutable that XML is getting used, it has yet to come even close to living up to the exalted expectations that were set as for its applicability and acceptance.

The new Office 2003, due out later this year, will feature integrated XML support across the entire application suite. That will certainly make XML more visible, accessible and real. But, XML support has been available with Excel for awhile now – as shown here.

Ironically, Web services is increasingly looking like the face-saving *killer application* for XML. But the lack of commercial 'success' has not in anyway deterred the intelligentsia from having a grand old time with XML. If you look around there are all sorts of interesting flavors of XML for specific scientific disciplines. For example, there is a [Chemical Markup Language \(CML\)](http://www.xml-cml.org) for representing molecular information [www.xml-cml.org]. This makes sense since chemical structures, especially molecular structures, are difficult to represent in a non-graphical, textual form.

Then we have the [Human Markup Language \(HumanML\)](http://www.humanmarkup.org) [www.humanmarkup.org] to enable consistent description of human emotions, intentions, gestures etc. That does

```
- <Worksheet ss:Name="Sheet1">
- <Table ss:ExpandedColumnCount="8" ss:ExpandedRowCount="34" x:FullColumns="1" x:FullRows="1">
  <Column ss:AutoFitWidth="0" ss:Width="106.5" />
  <Column ss:AutoFitWidth="0" ss:Width="46.5" />
  <Column ss:AutoFitWidth="0" ss:Width="43.5" />
  <Column ss:Hidden="1" ss:AutoFitWidth="0" />
  <Column ss:AutoFitWidth="0" ss:Width="39" />
  <Column ss:AutoFitWidth="0" ss:Width="106.5" />
  <Column ss:Index="8" ss:AutoFitWidth="0" ss:Width="49.5" />
- <Row ss:AutoFitHeight="0" ss:Height="27">
  - <Cell ss:MergeAcross="7" ss:StyleID="s27">
    <Data ss:Type="String">Anu's & Deanna's Wedding - January 1, 2003</Data>
  </Cell>
  </Row>
- <Row ss:Index="3">
  - <Cell ss:Index="6" ss:StyleID="s31">
    <Data ss:Type="String">Status as of =</Data>
  </Cell>
  - <Cell ss:MergeAcross="1" ss:StyleID="s44" ss:Formula="=DATE(2002,11,30)">
    <Data ss:Type="DateTime">2002-11-30T00:00:00.000</Data>
  </Cell>
  </Row>
- <Row ss:Index="5" ss:Height="25.5">
  - <Cell ss:StyleID="s39">
    <Data ss:Type="String">Maximum number of guests possible =</Data>
  </Cell>
  - <Cell>
    <Data ss:Type="Number">100</Data>
  </Cell>
  </Row>
- <Row ss:Index="7">
  - <Cell>
    <Data ss:Type="String">Total invited todate =</Data>
  </Cell>
  - <Cell ss:Formula="=R[23]C+R[23]C[5]">
    <Data ss:Type="Number">25</Data>
  </Cell>
  </Row>
- <Row ss:Index="9">
  - <Cell>
    <Data ss:Type="String">RSVP todate =</Data>
  </Cell>
  - <Cell ss:Formula="=R[21]C[1]+R[21]C[6]">
    <Data ss:Type="Number">10</Data>
  </Cell>
```

A very small snippet of Microsoft Excel's representation of a spreadsheet as an XML document. As is to be expected the XML representation is somewhat lengthy and convoluted.

raise an eyebrow – and no doubt there is schema within HumanML to express that!

But then there is the W3 sanctioned [Mathematical Modeling Language \(MathML\)](http://www.w3.org/TR/REC-MathML) [www.w3.org/TR/REC-MathML] for describing mathematical notation – and that makes you stop and think. Before we go any further lets see how you would describe a rather simple mathematical representation such as: $(a + b)^2$

In MathML this would be shown as – where **<mi>**, **<mn>** and **<mo>** represent identifiers, numbers and operands, respectively:

```
<msup>
  <mfenced>
    <mrow>
      <mi>a</mi>
```

```

    <mo>+</mo>
    <mi>b</mi>
  </mrow>
</mfenced>
<mn>2</mn>
</msup>

```

OK ... so far?

Let us try something a bit more complicated such as:

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

The MathML markup would look like:

```

<mrow>
  <mi>x</mi>
  <mo>=</mo>
  <mfraction>
    <mrow>
      <mrow>
        <mo>-</mo>
        <mi>b</mi>
      </mrow>
      <mo>&PlusMinus;</mo>
      <msqrt>
        <mrow>
          <msup>
            <mi>b</mi>
            <mn>2</mn>
          </msup>
          <mo>-</mo>
          <mrow>
            <mn>4</mn>

```

```

    <mo>&InvisibleTimes;</mo>
    <mi>a</mi>
  <mo>&InvisibleTimes;</mo>
  <mi>c</mi>
</mrow>
</mrow>
</msqrt>
</mrow>
<mrow>
  <mn>2</mn>
  <mo>&InvisibleTimes;</mo>
  <mi>a</mi>
</mrow>
</mfraction>
</mrow>

```

Now you have to ask whether this is really necessary. I can think of quite a few respected figures, Ken Iverson of IBM who created APL for a start, who will persuasively argue that mathematical notation is self-explanatory and unambiguous. They will contend, cogently, that XML does not come even close to being as universal and intuitive as mathematical notation.

Plus don't we have nearly four decades of experience with compiler technology for 'sorting out' mathematical representation – not to even mention even older schemes such as [Reverse Polish Notation](#). Given all of this do we really need a MathML to clearly and succinctly represent mathematical notations? I don't know. What do YOU think?

Knowledge Management and Corporate Portals

Last month we continued our ongoing series [Dec. 2002] on Knowledge Management (KM) vis-à-vis corporate portals. This concludes this series.

Corporate portals can facilitate both codified and noncodified KM, where codified refers to IT-based schemes with noncodified thus alluding to human-centric schemes. Portals include as standard functions many of the

core services necessary to implement successful KM applications – key among these being collaborative tools, wide ranging information dissemination features, powerful security for access control, data syndication capabilities and personalization.

Thus, as with other ERP applications, there can be significant synergy in integrating KM functionality with a corporate portal.

However, as with other ERP applications, KM and portals, in the end, are still separate initiatives and endeavors.

Though the glut of ready information available via the Web makes KM more compelling than ever before, a company's attitudes about KM is likely to predate its interest in implementing a corporate portal. If your company already has one or more KM applications, then as with ERP, the first thing that needs to be done in terms of the portal is to ensure that these KM applications can be accessed via the portal.

If the KM applications run on host systems or already have Web interfaces, then the portal integration should be relatively simple as with other similar host applications. However, some KM applications rely on a client-server model, with desktop processing power being used to do some of the data analysis, modeling and reporting. Adapting such client-server KM application for use via the portal could be more challenging.

In the case of company employees already using the client-server KM applications, it might actually just be simpler and logical to leave them as they are to begin with – and just treat them akin to desktop productivity applications [e.g. Word, Excel or PowerPoint]. The users will still be able to use the portal for some KM related collaborative, information publishing and data transfer functions, but the client-server applications *per se* will be invoked and executed outside the Web browser windows used to interact with the portal. If these KM applications are only used by a small number of users [e.g. less than 50], it would be difficult to justify the cost and effort of trying to implement a fully integrated solution.

The basic flaw of this approach is that these KM applications cannot be readily offered to external users [e.g. select partners] via the

portal. If external access makes business sense, the easiest solution might be to evaluate a Web-oriented 'terminal server' scheme such as Citrix's NFuse – which will allow the client component to execute on a server and deliver its results in a browser window. The other option would be to investigate if a Web-enabled version of the product is available.


Once you have portal accessible KM applications, the issues vis-à-vis these applications and the portal are really the same as those pertaining to ERP applications. One major difference, however, could be that company's are unlikely to be as amenable to sharing their KM applications with external users as they are with ERP applications. The tangible cost-reduction possibilities of e-business makes ERP 'sharing' hard to resist. KM, however, is different. It is not an e-business process *per se*. It can influence, shape and monitor e-business – but unlike with say SCM, partners and suppliers do not have to access your KM tools in order for you to gain the necessary knowledge. That is the rub when it comes to KM applications.

Knowledge is a valuable and leveragable asset. Letting select partners and associates have ready access to all the relevant KM tools might not be entirely prudent from a business standpoint. It might just be better and safer to share selected insights using collaborative tools, customized portal-view pages, teleconferences, or face-to-face meeting. But this yet again will be a line-of-business decision as opposed to a portal related issue. The same will be true about the use of KM within the company. Determining the best ways to exploit KM has to be a line-of-business decision. A portal can provide the infrastructure for KM applications but it should not be a determining factor as to how and why KM should be used within a company.

i-sight: i-interesting i-business i-info: now with Deanna Gurugé

Sun turns to 32-bit, low-end, low-cost servers:

The space race analogy used in the February issue to characterize the tussle between Sun and IBM for server market domination continues to be rather germane and prescient. As IBM relentlessly pushes the performance envelope at the high end, as with the new 1.7GHz POWER4+-based p690 [page 11], Sun, having already made a major commitment to blade-computing [Feb. 2003], is now turning its attention to the low-end.



The Intel Xeon powered Sun Fire V60x servers capable of running either Solaris 9 for x86 or Red Hat Linux.

source: Sun

On May 27th Sun introduced two new, 32-bit, x86 machines, capable of running Solaris 9 for x86 or Red Hat Linux; viz. Sun Fire V60x and Sun Fire V65x. These aggressively priced boxes, i.e. a start price of \$2,450 and \$2.650 respectively, rather than using Sun's UltraSPARC processors are instead powered by Intel Xeon processors.

The basic hardware characteristics of these two boxes are as follows:

- up to two 3.06-GHz Intel Xeon microprocessors (where these

'hyperthreadable', long pipeline processors are specifically designed for server-side applications)

- up to 12GB RAM
- up to six 100-MHz PCI-X [Jan. 2003] slots
- A 533-MHz front-side bus
- Dual Gigabit Ethernet

Sun is promoting Oracle software on these boxes – irrespective of whether they are running Solaris or Linux. The message and goal is here is very clear. With these low-end boxes Sun is going after the Windows 2000 segment of the market.

NetManage changes color like a chameleon?:

While it is safe to start writing-off NetManage as a force in the host access market, one has to admit that following its machinations can become addictively compelling given the sheer, unadulterated entertainment value. Though they are supposedly consolidating their marketing in Cupertino, CA [Feb. 2003], Bertram Rankin, who entertained us with so much panache just last February, is now toast. So is the V.P. R&D.

1Q2003 results, announced at the end of April were not good. Net revenue was \$14.1M, compared to \$19M in 1Q2002. Net loss for the quarter was \$1.2 million.



NetManage stock on the re-bounce despite systematically declining revenues.

source: Excite

But the stock after getting pretty close to \$1 has nearly doubled – mainly in the continued belief that there is 'cash' left in this stock. This really is better than any soap on TV.

TEL: (603) 293-5855

WWW.WOWNH.COM & WWW.INET-GURU.COM

E-MAIL: anu@wownh.com & deanna@wownh.com