Techline
Making it easier to deliver client value

IBM

Metropolitan VM Users Association Meeting
April 20, 2010

Liz Holland
System z Certified Consulting IT Specialist

Agenda

- **Speaker introduction**
- **System z explained to ISVs new to the platform**
  - From a presentation given at the Dallas IBM Innovation Center March 10, 2010
  - Audience - Participants in the Chiphopper program
- **zCP3000 enhancements to support IFLs and z/VM Capacity Planning**

Techline
Making it easier to deliver client value

IBM

Today's Speaker

Liz Holland celebrated 30 years with IBM on February 20, 2010

- **Started fixing Selectric typewriters**
- **Certified Consulting IT Specialist for System z**
- **Currently part of IBM Global Techline**
- **Capacity planning experience for MVS since 1988**
- **Linux on the mainframe support since 1998 (the beginning)**
  - Original Team Lead of the Chiphopper program
- **Author of numerous documents on z/VM Capacity Planning**
- **Jointly developing new features for zCP3000 to enhance z/VM planning**

(Notes:The Chiphopper program gives software developers free access to Linux on IBM platforms: Power, x and z. In the first year, about 230 new products for Linux on System z were developed. Currently the number is over 3,000.)

“Gloabl Techline” means that IBM has in place a unified approach to Technical pre-sales support. The offerings are available globally.

The Chiphopper program gives software developers free access to Linux on IBM platforms: Power, x and z. In the first year, about 230 new products for Linux on System z were developed. Currently the number is over 3,000.

- This presentation was given to an audience of Independent Software Vendors at the Dallas Innovation Center on March 10, 2010.

Mixed Workloads. System z strength - Consolidation Works.

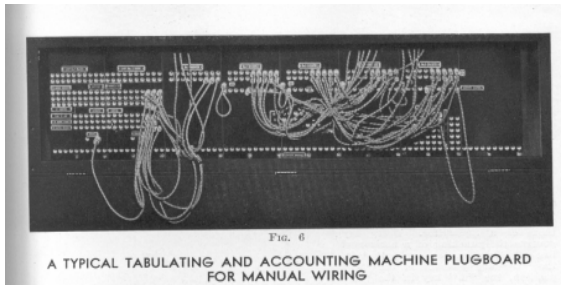

FIG. 6

A TYPICAL TABULATING AND ACCOUNTING MACHINE PLUGBOARD FOR MANUAL WIRING

- **An early IBM computer control board for the 407 tabulating machine.**
- **The Hardware Save Area in System z is the remaining vestige of this plugboard – the interface between the hardware and the software.**
- **The Hardware Management Console provides control and access to I/O.**

The program is changed by moving the patch cords.

Another patch cord technology that virtualizes well is the phone company – look at your phone, with the internet on it.

Mixed Workloads:System z strength and Why Consolidation Works.

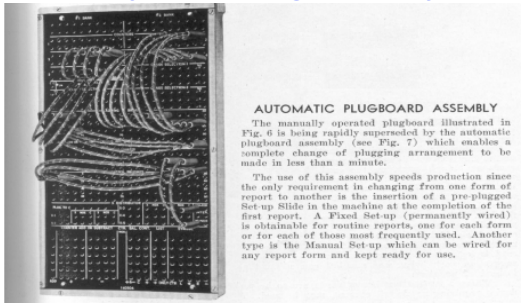

AUTOMATIC PLUGBOARD ASSEMBLY

The manually operated plugboard illustrated in Fig. 6 is being rapidly superseded by the automatic plugboard assembly (see Fig. 7) which enables a complete change of plugging arrangement to be made in less than a minute.

The use of this assembly speeds production since the only requirement in changing from one form of report to another is the insertion of a pre-plugged Set-up Slide in the machine at the completion of the first report. A Fixed Set-up (permanently wired) is obtainable for routine reports, one for each form or for each of those most frequently used. Another type is the Manual Set-up which can be wired for any report form and kept ready for use.

- **A more advanced pre-wired assembly.**
- **Think of 'task switching' as removing and inserting a complete assembly – complete program isolation with EAL5 security.**

Is that 'assembly' where Assembler language came from?

Speaker notes for Slides 5 and 6

Slide 5. The program is changed by moving the patch cords.

Another patch cord technology that virtualizes well is the phone company – look at your cell phone, with the internet on it.

Slide 6. Is that 'assembly' where the term Assembler Language came from?

System Design Affects Virtualization Capabilities

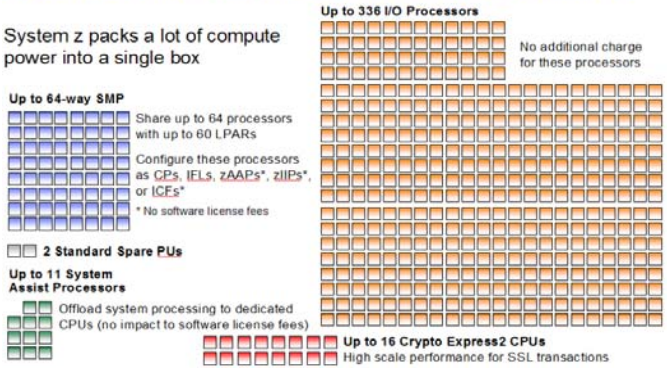

A 64-way System z actually has many more ancillary processors available.
These do not incur software license charges.
Permits consolidating 30 distributed processors(cores) to 1 System z processor (IFL)!

System Design Affects Virtualization Capabilities

Compare to typical UNIX system design...

Up to 128-way SMP configuration

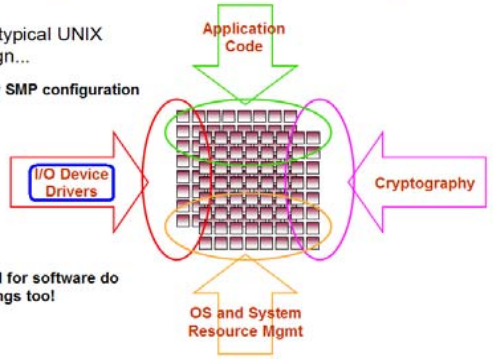

CPUs licensed for software do a lot other things too!

Why does I/O offload matter so much?
If a computer chip operated in seconds (instead of GHz), it would take 11 days for an I/O operation to complete. How many other things can you do in 11 days?

Techline
Making it easier to deliver client value

IBM

What Makes A Best Fit Workload for Linux on System z?

- **Leverage classic strengths of the System z**
  - High availability
  - High i/o bandwidth capabilities
  - Flexibility to run disparate workloads concurrently
  - Requirement for excellent disaster recovery capabilities
  - Security
- **Shortening end to end path length for applications**
  - Co-location of applications
  - Consolidation of applications from distributed servers
  - Reduction in network traffic
  - Simplification of support model
- **Consolidation Effect**
  - **Power requirements**
  - **Software costs**
  - **People Costs**
  - **Real Estate**
  - **Workloads requiring EXTREME Flexibility**

Framing the discussion around 'Fit for Purpose'. Choose the right deployment model for the application.

Techline
Making it easier to deliver client value

System z Achieves High Core-to-Core Ratios When Consolidating from Distributed Environments

- Demonstrates unique IBM Value Proposition against competitive scale-out solutions
- **Real customers, real workloads!**

Customer	Distributed Cores	IBM System z10™ or IBM System z9® Cores	Ratio of distributed to System z cores*
Nationwide	450	21 cores z9	21 to 1
Pension Fund Agency	1324	36 cores z10	36 to 1
Government Agency	292	5 cores z10	58 to 1

Contact IBM to help with your specific situation!

* Client results will vary based on each specific customer environment including types of workloads, utilization levels, target consolidation hardware, and other implementation requirements.

Background information which illustrates some consolidation ratios. Actual results from customers' production environments.

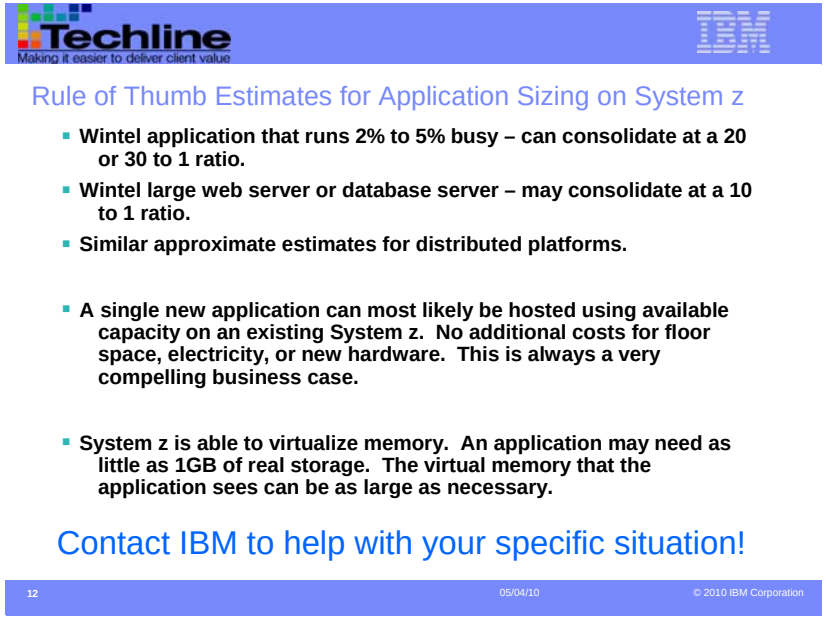

Very general guidance to ISVs considering putting their applications on System z. To be used as a starting point. Benchmarking is always preferable (but expensive).

Speaker notes for Slides 10, 11 and 12

Slide 10. Framing the discussion around 'Fit for Purpose'. Choose the right deployment model for the application.

Slide 11. Background information which illustrates some consolidation ratios. Actual results from customers' production environments.

Slide 13. Very general guidance to ISVs considering putting their applications on System z. To be used as a starting point. Benchmarking is always preferable (but expensive).

- This presentation was given to an audience of Independent Software Vendors at the Dallas Innovation Center on March 10, 2010.

IFLs and z/VM

- **New Support to add an IFL only CEC (Jan 2010)**
  - Several step process
- **Additional valid workload mixes**
  - Agreed to by z/VM Lab, ATS, and Techline
  - Based on z/OS measurements
- **We believe all the graphs now refer to IFLs when they should**
- **The number of IFLs is presented in the Smart Text**
  - There is no comprehensive designator for machines with GPs and IFLs
  - A 2097-706 E18 could have from zero to 12 IFLs.
    - Same ambiguity for zIIPs and zAAPs

Graphs correctly refer to IFLs

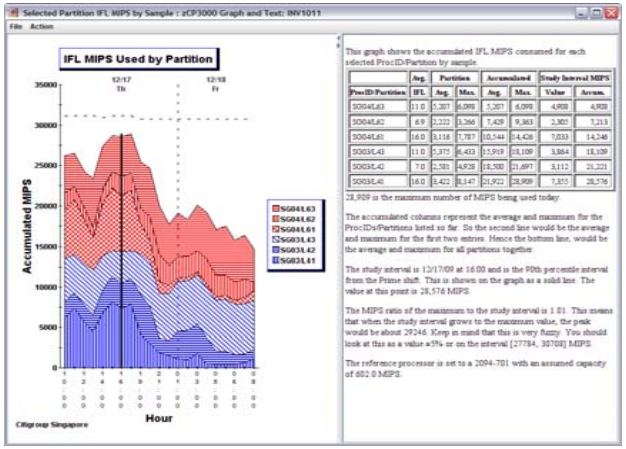

IFLs and other engine types called out in Smart Text

The Study Interval (the 90th percentile interval from the Prime shift) is shown on the graph as a solid line drawn on 12/17/09 at 16:00.

	Study	Max.	LPAR	PUs		Capacity		Workload
Partition	CPU%	CPU%	Weight%	#	Type	Min	Max	Mix
L41	35.0%	38.7%	64%	16.0	IFL	9,615	9,615	TM-Mix/LV
C41	2.8%	2.9%	Ded	1.0	ICF	653	653	CFCC
C42	2.7%	2.8%	85%	1.0	ICF	546	643	CFCC
C43	2.7%	2.9%	Ded	1.0	ICF	653	653	CFCC
C44	0.1%	0.1%	15%	1.0	ICF	96	643	CFCC
P41	10.2%	13.4%	63%	6.0	CP	2,562	4,058	LoIO-Mix
P41 zIIP	1.1%	1.1%	71%	1.0	zIIP	489	694	
P44	0.7%	1.2%	8%	2.0	CP	348	348	LoIO-Mix
T41	5.0%	5.7%	26%	4.0	CP	1,079	2,727	LoIO-Mix
T41 zIIP	0.2%	0.2%	29%	1.0	zIIP	210	714	
T44	0.3%	0.4%	2%	2.0	CP	99	99	LoIO-Mix
L42	13.4%	22.0%	20%	5.0	IFL	3,325	3,325	TM-Mix/LV
L43	18.7%	31.3%	16%	11.0	IFL	2,360	6,491	TI-Mix/LV

Linux guests

- **Virtual memory information**
  - Amount defined in the directory
  - Working set size (max; min; at time of study interval)
- **Will be adding # of virtual CPUs**
  - Requires calculation which will be done in new version of the extract program

Memory Summary

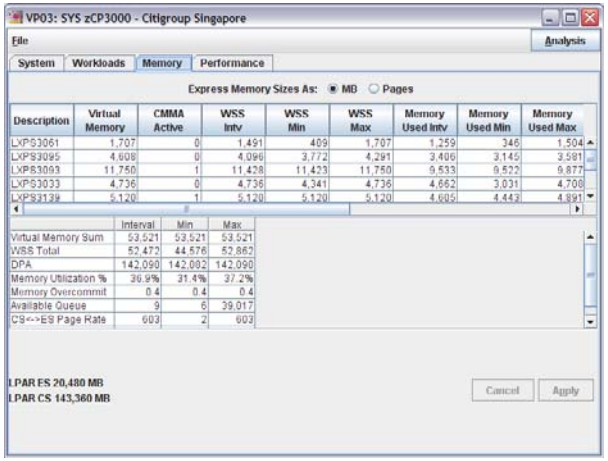

Call to Action

- **Collect more data**
  - Techline capacity planning and sizing teams are ready to help
  - z/VM Capacity Planning data collection
- **Build a business case to justify more enhancements**
  - LSPR measurements of more z/VM workloads
  - Implement CPU Measurement Facility in z/VM
- **Help define a better Health Check**
  - What key indicators are important to you

Useful Links

- **Executive Overview of zCP3000 (for z/VM customers)**
  - http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS3948
- **z/VM Capacity Planning Analysis – Data Collection**
  - http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS2875