Tämän päivän mainframe
LSU 8-9.11.2005

Tapio Koskinen
IBM Systems and Technology Group

40 years of Mainframes

Celebrating the 40th anniversary of the S/360 mainframe.

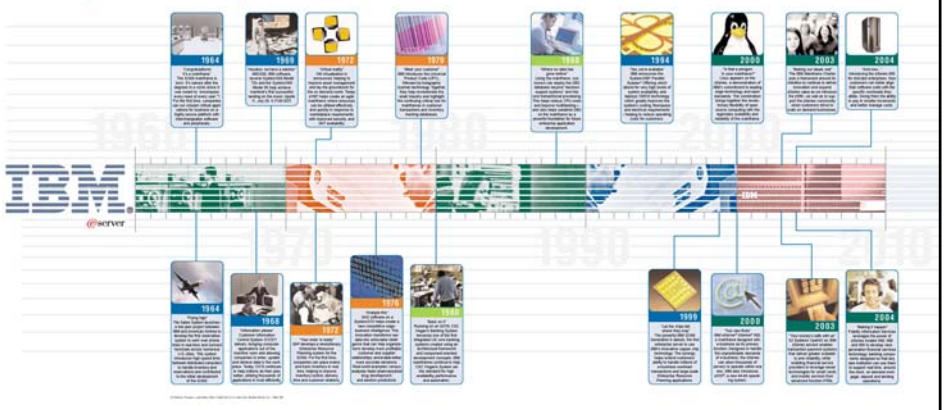

System/360 Model 40

0,03 – 0,75 MIPS
8KB – 8MB Memory
~ 3 M$/mip

26 – 1365 MIPS
8GB – 32GB Memory

z890

IBM System z9 and eServer zSeries

IBM eServer zSeries 900 z900 (2064)

- Announced 10/00 - first 64-bit zSeries
- 42 models – Up to 16-way
- Specialty Engines
  - CP, IFL, ICF
- On Demand Capabilities
  - CUoD, CIU, CBU
- Memory – up to 64 GB
- Channels
  - Up to 256 ESCON® channels
  - FICON® Express, Parallel
  - Token- Ring, FDDI, Ethernet, ATM
  - Coupling Links
- Crypto coprocessors, accelerators
- Parallel Sysplex clustering
- HiperSockets™ – up to 4
- Up to 15 logical partitions
- Operating Systems
  - OS/390®, z/OS®, VM/ESA®, z/VM®, VSE/ESA™, z/VSE™, TPF/ESA, z/TPF, Linux® on zSeries

IBM eServer zSeries 800 z800 (2066)

- Announced 2/02 - first 64-bit zSeries for mid market
- 10 models – Up to 4-way
- Specialty Engines
  - CP, IFL, ICF
- On Demand Capabilities
  - CUoD, CIU, CBU
- Memory – up to 32 GB
- Channel
  - Up to 240 ESCON Channels
  - FICON Express
  - Networking Adapters (OSA)
  - Coupling Links
- Cryptographic Coprocessors
- Parallel Sysplex clustering
- HiperSockets – up to 4
- Up to 15 partitions
- Operating Systems
  - OS/390, z/OS, VM/ESA, z/VM, VSE/ESA, z/VSE, TPF/ESA, z/TPF, Linux on zSeries

IBM eServer zSeries 990 z990 (2084)

- Announced 5/03 – First zSeries Superscalar Server
- 4 models – Up to 32-way
- Specialty Engines
  - CP, IFL, ICF, zAAP
- On Demand Capabilities
  - CUoD, CIU, CBU, On/Off CoD
- Memory – up to 256 GB
- Channels
  - Four LCSSs
  - Up to 1024 ESCON channels
  - Up to 240 FICON Express2 channels
  - Token-Ring, GbE, 1000BASE-T Ethernet
  - Coupling Links
- Crypto Express2
- Parallel Sysplex clustering
- HiperSockets - up to 16
- Up to 30 logical partitions
- Operating Systems
  - OS/390, z/OS, z/VM, VSE/ESA, z/VSE, TPF/ESA, z/TPF, Linux on zSeries

IBM eServer zSeries 890 z890 (2086)

- Announced 4/04 – zSeries Superscalar Server for mid market
- 1 model – Up to 4-way
  - 28 capacity settings
- Specialty Engines
  - CP, IFL, ICF, zAAP
- On Demand Capabilities
  - CUoD, CIU, CBU, On/Off CoD
- Memory – up to 32 GB
- Channel
  - Two LCSSs
  - Up to 420 ESCON channels
  - Up to 80 FICON Express2 channels
  - Networking Adapters (OSA)
  - Coupling Links
- Cryptographic Coprocessors
- Parallel Sysplex clustering
- HiperSockets - up to 16
- Up to 30 partitions
- Operating Systems
  - z/OS, z/VM, VSE/ESA, z/VSE, TPF/ESA, z/TPF, Linux on zSeries

IBM System z9 109 z9-109 (2094)

- Announced 7/05
- Superscalar Server
- 5 models – Up to 54-way
- Specialty Engines
  - CP, IFL, ICF, zAAP
- On Demand Capabilities
  - CUoD, CIU, CBU, On/Off CoD
- Memory – up to 512 GB
- Channels
  - Four LCSSs
  - Multiple Subchannel Sets
  - MIDAW facility
  - 63.75 subchannels
  - Up to 1024 ESCON channels
  - Up to 336 FICON Express2 channels
  - 10 GbE, GbE, 1000BASE-T
  - Coupling Links
- Configurable Crypto Express2
- Parallel Sysplex clustering
- HiperSockets - up to 16
- Up to 60 partitions
- Enhanced Availability
- Operating Systems
  - z/OS, z/VM, VSE/ESA, z/VSE, TPF/ESA, z/TPF, Linux on System z9

Mainframes Continues to Scale with z9-109

- Each new range is designed to deliver:
  - New function
  - Improved performance
  - Improved availability and serviceability
  - Field upgrades to help protect customer investment
  - Better price/performance

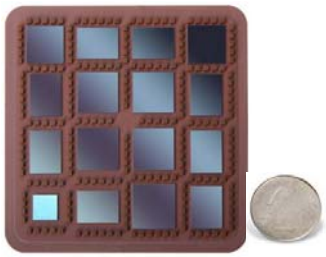

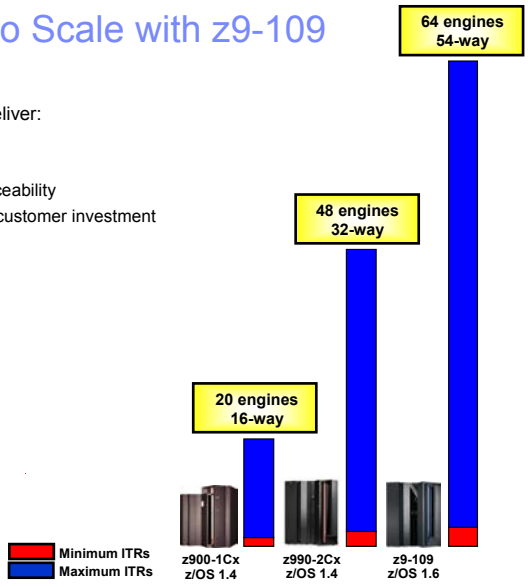

System z9 and zSeries: Balanced System Design

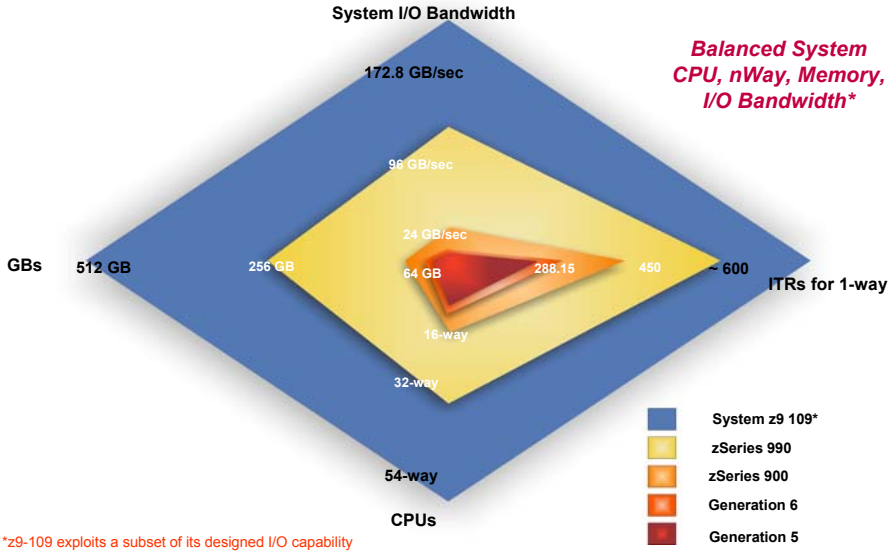

*z9-109 exploits a subset of its designed I/O capability

IBM System z9 109 Overview

- Machine Type
  - 2094
- 5 Models
  - S08, S18, S28, S38 and S54
- Processor Units (PUs)
  - 12 (16 for Model S54) PUs per book
  - 2 SAPs per book, standard
  - 2 spares per server
  - 8, 18, 28, 38 or 54 PUs available
    - Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z9 Application Assist Processors (zAAPs), optional System Assist Processors (SAPs)
- Memory
  - Minimum of 16 GB
  - Up to 128 GB per book
  - Up to 512 GB for System
    - 16 GB increments
- I/O
  - Up to 16 STIs per book @ 2.7 GB/s each
  - Up to 4 Logical Channel Subsystems (LCSSs)
  - Total system I/O bandwidth capability of 172.8 GB*

* z9-109 exploits a subset of its designed I/O capability

z9-109 New Functions and Features

- Five new hardware models
- Faster Uni Processor
- Up to 54 CPs
- Up to 512 GB memory
- Up to 60 LPARs
- CBU for IFL, ICF and zAAP
- Separate PU pool management
- Redundant I/O interconnect
- Enhanced Driver maintenance
- Enhanced Book availability
- Dynamic oscillator switchover
- 54 additional hardware instructions

- Hot pluggable/ maintainable MBA/STI fanout cards
- Up to 16 2.7 GB STIs per book
- MIDAW facility
- Multiple Subchannel Sets per LCSS
- 63.75K Subchannels for Set-0
- Increased Number of FICON Express2 Features
- N_Port ID Virtualization
- IPv6 Support for HiperSockets
- OSA-Express2 1000BASE-T
- OSA-Express2 OSN (OSA for NCP)
- Enhanced CPACF with AES, PRNG and SHA-256
- Configurable Crypto Express2

*This statement represents IBM's current intentions. IBM development plans are subject to change or withdrawal without further notice.

z9-109 – Under the covers (Model S38 or S54)

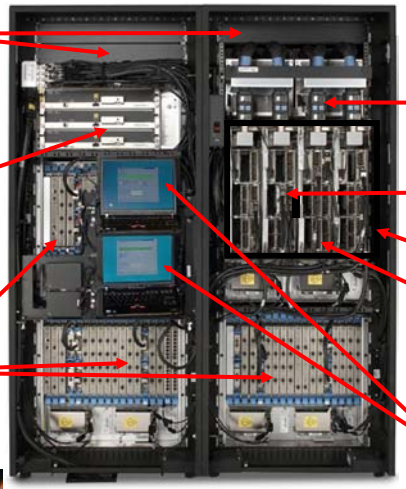

Front View

z9-109 12-way MCM

- Advanced 95mm x 95mm MCM
  - 104 Glass Ceramic layers
  - 16 chip sites, 217 capacitors
  - 0.476 km of internal wire

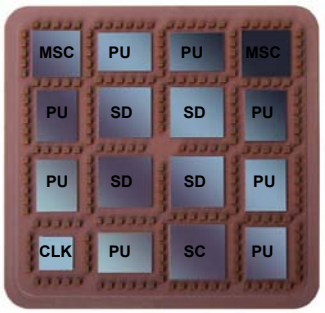

- CMOS 10K chip Technology
- PU, SC, SD and MSC chips
- Copper interconnections, 10 copper layers
- 8 PU chips/MCM
  - 15.78 mm x 11.84 mm
  - 121 million transistors/chip
  - L1 cache/PU
    - 256 KB I-cache
    - 256 KB D-cache
  - 0.58 ns Cycle Time
- 4 System Data (SD) cache chips/MCM
  - 15.66 mm x 15.40mm
  - L2 cache per Book
    - 660 million transistors/chip
    - 40 MB
- One Storage Control (SC) chip
  - 16.41mm x 16.41mm
  - 162 million transistors
  - L2 cache crosspoint switch
  - L2 access rings to/from other MCMs
- Two Memory Storage Control (MSC) chips
  - 14.31 mm x 14.31 mm
  - 24 million transistors/chip
  - Memory cards (L3) interface to L2
  - L2 access to/from MBAs (off MCM)
- One Clock (CLK) chip - CMOS 8S
  - Clock and ETR Receiver

z9-109 Processor Book Layout

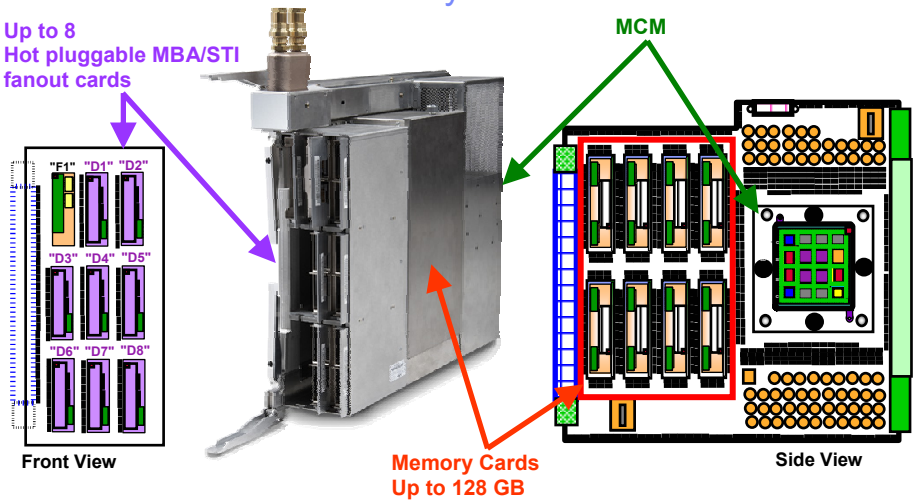

Note: 1. Concept Illustration only - not to scale
2. 4 or 8 pluggable Memory Cards
3. Each MBA fanout card is hot-pluggable and has 2 STIs

z9-109 Model S38 – Communication Ring Structure

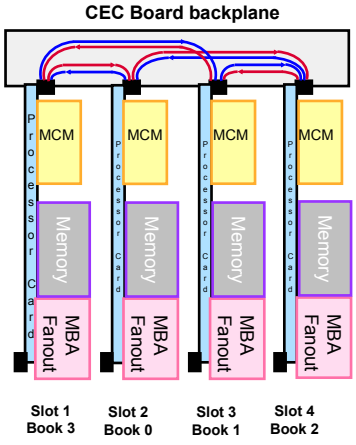

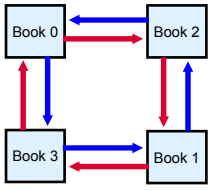

- **The ring structure consists of two rings (one running clockwise, the other counterclockwise)**
- **In a two or three Book configuration, Jumper Book(s) (installed in the CEC cage) complete the ring**
  - Jumper Books are not needed for a single-Book configuration
- **Books may be able to be inserted into or removed from the ring nondisruptively***
  - May allow **Concurrent book add** for model upgrade
  - **Enhanced book availability** to return a book after removal for upgrade or repair

* Customer pre-planning required, may require acquisition of additional hardware resources

z9-109 PU Characterization

- The type of Processor Units (PUs) that can be ordered on z9-109:
  - Central Processors (CPs)
    - Provides processing capacity for z/Architecture™ and ESA/390 instruction sets
    - Runs z/OS, z/VM, VSE/ESA, z/VSE, TPF/ESA, z/TPF, Linux for System z9 and zSeries and Linux under z/VM or Coupling Facility
    - z9-109 has Capacity Marker features NOT Unassigned CP features
  - IBM System z9 Application Assist Processors (zAAPs)
    - Under z/OS, the Java™ Virtual Machine (JVM) assists with Java processing to a zAAP
  - Integrated Facility for Linux (IFL)
    - Provides additional processing capacity for Linux workloads
  - Internal Coupling Facility (ICF)
    - Provides additional processing capacity for the execution of the Coupling Facility Control Code (CFCC) in a CF LPAR
  - System Assist Processors (SAPs)
    - SAPs manage the start and ending of I/O operations for all Logical Partitions and all attached I/O

Requirements for zAAP Exploitation

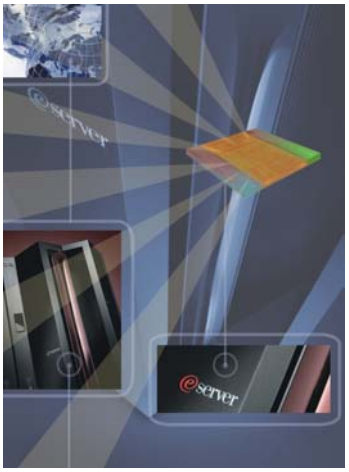

- Available on z9-109, z990, and z890
- Prerequisites:
  - z/OS 1.6 (or z/OS.e 1.6 on z890)
  - IBM SDK for z/OS, Java 2 Technology Edition, V1.4 with PTF for APAR PQ86689
- Subsystems and Applications using SDK 1.4 will exploit zAAPs automatically:
  - WAS 5.1
  - CICS® /TS 2.3
  - DB2 V8
  - IMS™ V8
  - WebSphere WBI for z/OS
- zAAPs must be jointly configured with general purpose processors within z/OS LPARs
  - Number of zAAPs may not exceed the number of permanently purchased CPs (including z990 unassigned CPs or z890 Downgrade - Record Only CPs) on a given machine model.

z9-109 Performance Comparison

Number of z9-109 CPs	Base	Ratio
1 CP	z990 301	1.35
8 CPs	z990 308	1.34
18 CPs	z990 318	1.34
28 CPs	z990 328	1.32
38 CPs	z990 332	1.51
54 CPs	z990 332	1.95

Note: For MSU values, refer to:
www-1.**ibm.com**/servers/eserver/zseries/library/swpriceinfo/
For ITRs refer to: www-1.**ibm.com**/servers/eserver/zseries/lspr/zSerieszOS.html

* CI = Capacity Indicator and refers to number of installed CPs. Reported by STSI instruction. Model 700 does not have any CPs.

z890 – Overview

- **Memory**
  - 8 GB Standard
  - 8 GB increments to 32 GB (8, 16, 24, 32 GB)
  - One concurrent memory upgrade path (24 – 32 GB)
- **Support for up to 30 LPARS**
  - Except for Capacity setting 110 which supports 15 LPARS
- **Cryptographic coprocessor optional**
- **New packaging for I/O with Two Logical Channel SubSystems (LCSS)**
  - 28 slot I/O cage supports up to 420 ESCON® channels
    - z890 110 capacity setting only has 16 I/O slots available
  - OSA-Express – Gigabit Ethernet, 1000BASE-T Ethernet, Token-Ring, Integrated Console Controller
  - Open FCP
  - Quadrupled HiperSockets™ support over z800 to 16
- **Single frame**
  - One and three phase options
  - Raised floor recommended but not required
  - Internal Battery Option

z890 on demand – A new way to think about granularity

- **Single Machine : 2086 and a single Model : A04**
- **A dramatic new way to consider upgrading**
- **One MCM per model with 5 Processor Units (PUs)**
  - Four PUs available for characterization
    - CPs, Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), or zSeries Application Assist Processor (zAAPs)
  - One PU standard as an SAP
- **Standard CPs –**
  - Four full capacity processors **each** with 7 capacity settings
    - Entry point is approximately 32% less capacity than z800-0E1 and largest capacity setting is up to 123% more than z800-004
  - **Upgrades can be horizontal, vertical, diagonal**, to best fit your needs *

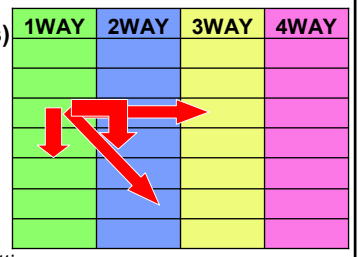

* Note: No mixing of standard CP capacity sizes in multi-engine machines, zAAPs cannot outnumber standard CPs in any machine.

Think of the possibilities:
Define the processor the way your business requires !

z9-109 Channel Type and Crypto Overview

- FICON/FCP
  - FICON Express
  - FICON Express2
- Networking
  - OSA-Express2
    - Gigabit Ethernet LX and SX
    - 10 Gigabit Ethernet LR
    - 1000BASE-T Ethernet
  - OSA-Express
    - Gigabit Ethernet LX and SX
    - 1000BASE-T Ethernet
    - Fast Ethernet (carry forward on upgrade)
  - HiperSockets
- Coupling Links
  - ISC-3 (Peer mode only)
  - ICB-3, ICB-4
  - IC
- ESCON
- Crypto
  - Crypto Express2
    - Configurable Coprocessor or Accelerator
- Channel types not supported:
  - OSA-Express Token-Ring (SOD Oct 2004)
  - PCIXCC
  - PCICA
  - ICB-2 (SOD 2003)
  - ISC-3 Links in Compatibility Mode (SOD April 2004)

Note: Only ICB cables orderable.
All other cables have to be sourced separately.

System z9 and zSeries Cryptographic Technology

- Continue to provide flexible Secure Sockets Layer (SSL) acceleration
- Continue to provide competitive symmetric performance in a security-rich environment
- Provide integration of Crypto features via ICSF
- Focus on required certifications and open standards
- Continue to improve performance
  - Each Crypto Express2 feature on a System z9, with both adapters configured as accelerators is designed to provide up to 6000* SSL handshakes per second

z900/z800 – Dec. 2000/ May 2002
2 Chips on CEC Board -
CMOS7s+ PCICC/PCICA (10/01)

G6 – June 1999
2 Chips on Processor
MCM - CMOS5x +
PCICC (6/99)

G5 – Sept. 1998
2 Chips on Processor
MCM - CMOS5x +
PCICC (6/99)

G4 – Sept. 1997
SCMs on Planar
Board - CMOS5x

G3 – June, 1997
SCMs on Planar
Board - CMOS5x

z9-109 – Planned for Sept, 2005
Crypto Express2

z990/z890 – January 2005
Crypto Express2

z890 – May 2004
PCIXCC/PCICA

z990 - September 2003
PCIXCC

z990 - June 2003
CPACF/PCICA

z900/z800 – Dec. 2000/ May 2002
2 Chips on CEC Board -
CMOS7s+ PCICC/PCICA (10/01)

*These measurements are examples of the maximum transactions/second achieved in a lab environment with no other processing occurring and do not represent actual field measurements. Details available upon request.

System z9 and zSeries Crypto Roadmap

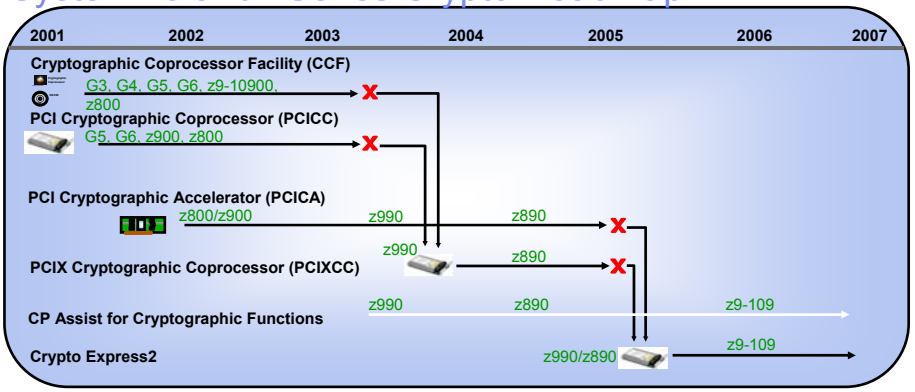

- Cryptographic Coprocessor Facility – Supports "Secure key" cryptographic processing
- PCICC Feature – Supports "Secure key" cryptographic processing
- PCICA Feature – Supports "Clear key" SSL acceleration
- PCIXCC Feature – Supports "Secure key" cryptographic processing
- CP Assist for Cryptographic Function allows "Clear key" crypto functions from any CP/IFL
- Crypto Express2 – Combines function and performance of PCICA and PCICC

System z9 and zSeries Availability and Scalability

- Parallel Sysplex
  - Availability / Capacity / Scalability
  - High level of performance
  - Internal or standalone CF
  - System-Managed CF Structure Duplexing
  - Designed for no single point of failure

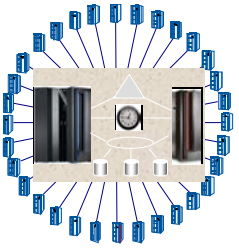

- Single System
  - Capacity Upgrade on Demand*
  - Capacity BackUp* extended to ICF, IFL, zAAP
  - Hot Pluggable I/O
  - Built in Redundancy
  - Advanced Problem Determination
    - FICON/FCP enhancements
  - Concurrent LIC Updates
  - Enhanced driver maintenance
  - Enhanced book availability
  - Redundant I/O Interconnect

System z9 continues to address the requirements for advanced availability and clustering

In Summary

- z9-109 provides a significant advancement in capabilities over its predecessors
  - Up to 35% faster uni processor performance
  - 54 new Hardware instructions
  - Up to twice the amount of memory
  - 80% increase in internal bandwidth
    - More STIs and more bandwidth per STI
  - Increased I/O capacity and constraint relief
  - On Demand capabilities
    - Significant advancements in unplanned and planned outage reduction
    - Virtualization advancements in LPAR and FCP N_Port ID virtualization