

FLEX-ES: A Bottom Up View of High Tech Mainframe Options

zFrame: a technical overview for



Mike Hammock
Cornerstone Systems Inc.
IBM zSeries Enablement Solutions
mhammock@csihome.com



Cornerstone's zFrame Objectives:

- Integrate proven software and hardware technologies using IBM based systems to support S/390 customers running OS/390, zOS, zVM, and VSE in the below 100 MIPS range.
- Provide a packaged solution capable of supporting mission critical production loads.
- Offer everything, including maintenance and services, at affordable price points.

What is "Cornerstone"

- ✓ CSI was established in 1991 as IBM's first business partner with a focus on system services.
- ✓ Became an IBM Business Partner Reseller in 1992
- ✓ CSI is currently one of IBM's largest Premier Business Partners
- ✓ HQ in Irvine, CA
- ✓ Offices located in Atlanta, Dallas, Valencia, Toronto, and London with coverage in Europe, Asia, and South America
- ✓ Focus on IBM zSeries Systems and Support including:
 - ✓ zOS, zVM, VSE, Websphere, DB2, Linux and UNIX porting, Lotus, and other zSeries products
- ✓ CSI is the largest reseller of FLEX-ES(tm) based systems - zFrame(tm) and the most experienced.



How zFrames are being used

- **Mission Critical Production**
 - **Insurance, Banking, Government, Military, Retail, Distribution**
- **Disaster Recovery or Backup Systems**
- **Out-sourcing**
- **In-sourcing**
- **Development and Test**



Hardware and Software Emulation - . . .some good, some bad, but all are part of our history . . .

- **1401 Emulation on S/360**
- **7090 and 7094 on S/360**
- **FBA**
- **VTape**
- **FakeTape**
- **P/370 and P/390**
- **CP/67**
- **VM**
- **VMTool**
- **Microcode**
- **LPAR**
- **First S/360 instructions were executed under emulation on a 7030 in 1963 . . .**



FLEX-ES Architectural Compliance

- **Meets all IBM definitions for the zSeries architecture (through ALS3)**
 - **Subjected to (and passed) the same verification tests as new IBM hardware**
- **Fully compatible with all current IBM operating systems and products**
- **Recommended by IBM for developers of new vendor software products**

This is a true, full function zSeries system, not just a highly capable API.



IBM CMOS - FLEX-ES similarities

- **IBM Systems**
- **FLEX-ES**

CP and SAP (System Assist Processor)	Emulated (CP) Processor and Host (SAP) Processor
CP Instruction Cache	Translated Instruction Cache
IOCP Configuration	<i>FLEX-ES Configuration File</i>
LPARs	Multiple Instances



Processor Performance



- **Currently available Intel Pentium technology allows (approximate):**
 - Single processor systems up to 40 MIPS
 - Dual processor systems up to 70 MIPS
 - Three processor systems up to 95 MIPS
- **S/370, ESA, or 64 bit zSeries mode**
 - 64 bit currently restricted to PWD systems
 - Runs z/OS 1.6 and zVM 5.1
- **3.5 GB of memory for S/390**
 - Much more memory is possible, but configuration restrictions limit usefulness



Current zFrame Models



Standard model base configurations

	MIPS (approx)	S/390 Memory	Disk (GB)	# of Proc.	xSeries
z8	8 (throttled)	3.4 GB	330 GB	2/1	232 2 x 1.4 GHz
z18	14-18	3.4 GB	330 GB	2/1	232 2 x 1.0 GHz
z30	32 - 36	3.4 GB	330 GB	2/1	235 2 x 3.2 GHz
z60	55 - 65	3.4 GB	330 GB	2/2	235 2 x 3.2 GHz
z90	90 - 100	3.4 GB	330 GB	4/3	255 4 x 3.0 GHz



IBM PWD Program

- PWD = (IBM's) PartnerWorld for Developers (ISVs)
- For Developers of S/390 software to be sold
- Special deals on IBM Software
 - "On loan" at no cost
- Special deals on *FLEX-ES* based solutions
- Special line of zFrame systems for these PWD members



Current zDev Models



- Standard zDev base configurations

	MIPS (approx)	S/390 Memory	Disk (GB)	# of Proc.	xSeries
zPad	36 (I/O Limited)	1.25 GB	60 GB	1/1	TP T42
zDev1	40	3.4 GB	160 GB	2/1	236 2 x 3.6 GHz
zDev2	70	3.4 GB	330 GB	2/2	236 2 x 3.6 GHz
zDev3	90 - 100	3.4 GB	330 GB	4/3	255 4 x 3.0 GHz



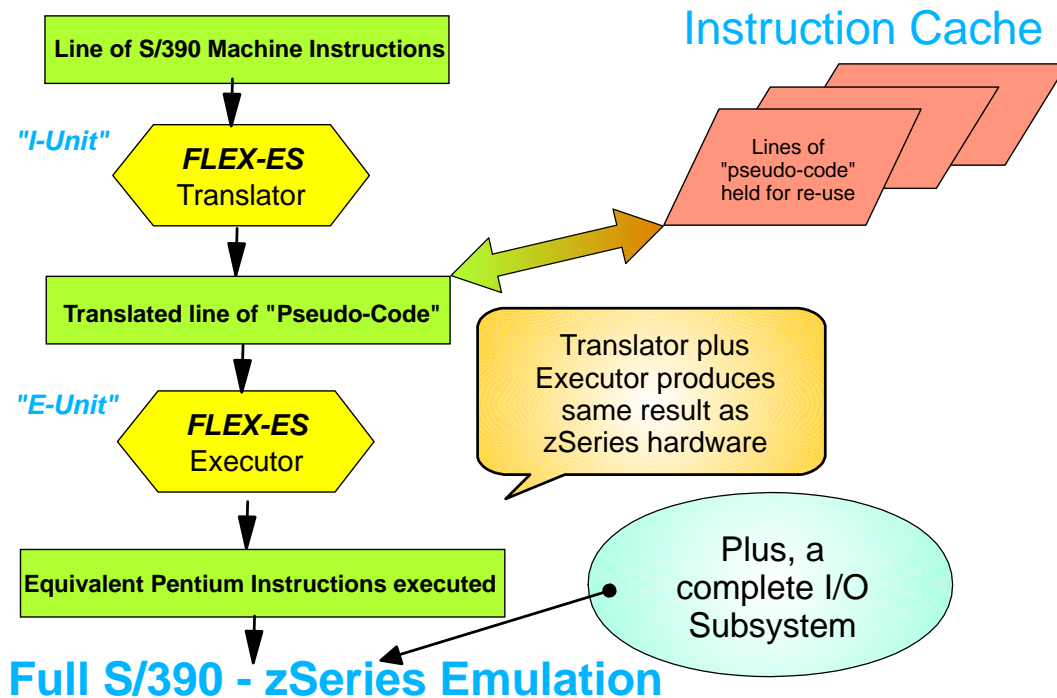
IBM Software License Charges

MIPS	License Base	Comments
8	ESL (OTC)	Inexpensive, OTC: (N/A: zOS, zVM 4.x)
16	1/3 GOLC	Attractive pricing. Limited to Pentium III 1.0 GHz
34	2/3 GOLC	Fastest Uni-Processor Less cost than 36 MIPS z800
60	GOLC/ZELC _(0A1)	Same license as MP3000 H30
95	ZELC (0B1)	Same license as 110 MIPS z800

License base is determined by IBM and is used on an "RPQ" basis uniquely for each system.

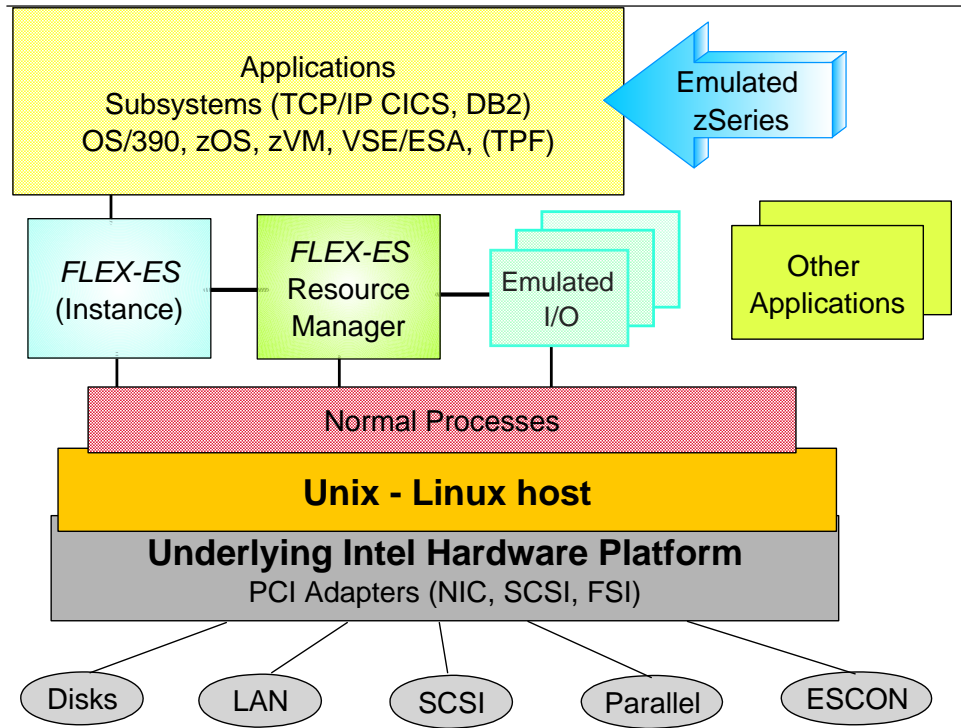


How does it work?

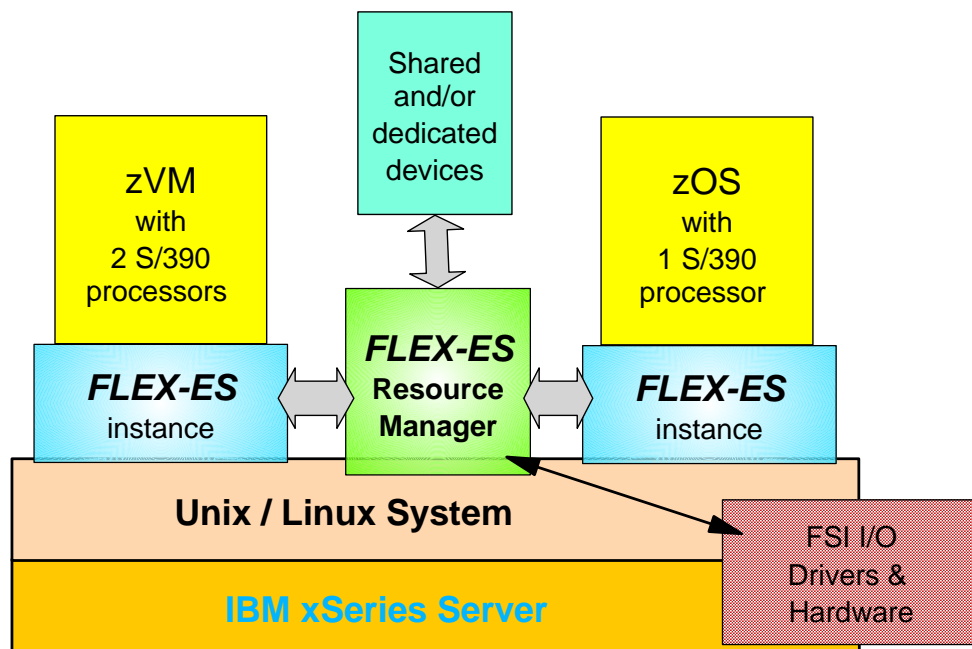




System Structure



Multiple Instances





How about running *TPF* ??

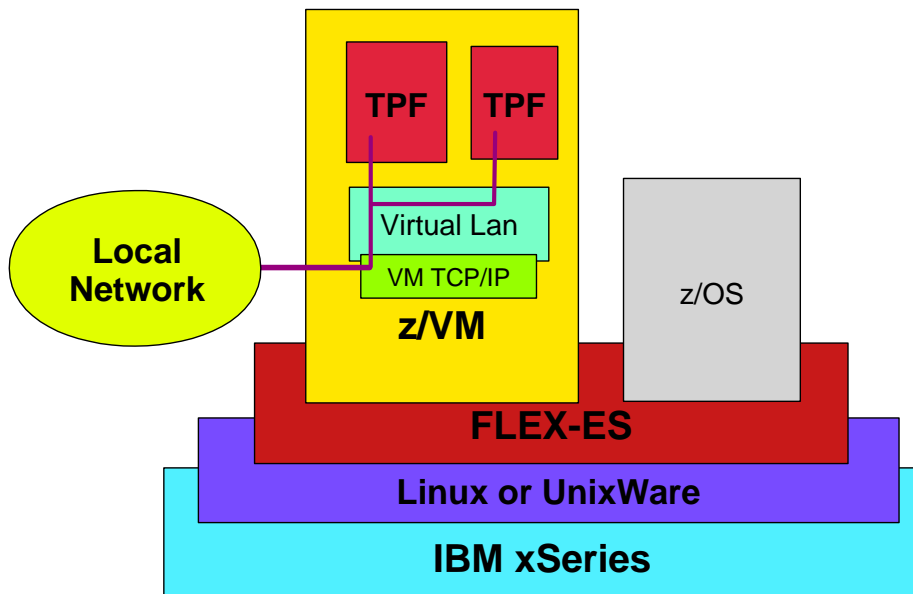
- **FLEX-ES does not support QDIO**
- **Current TPF requires QDIO for LAN communications**
- **Solution:**
 - ▶ **Run TPF as a zVM guest**
 - zVM runs very well on a zFrame
 - ▶ **Define a zVM Virtual LAN, which provides QDIO interface to TPF**
 - ▶ **Route LAN traffic through zVM to TPF**

Running as a zVM guest, TPF runs very well on the zFrame.

We have a number of TPF users and developers running on zFrame systems.



What it looks like





Disk Subsystem

- **Disk Capacity**
 - **1.4 or 2.8 TB internal disk capacity**
 - Effective S/390 space, in RAID-5 configuration
 - **Multiple more TB when using external disk enclosures**
 - **Very economical cost point**
- **Disk Performance**
 - **Very good performance: 1K - 3K IO / sec**
 - **See** http://www.perfassoc.com/jsc/pdf/papers/flex-es_io_performance_02.pdf
- **RAID-5 Disk array (IBM ServeRAID adapter)**
 - Full data redundancy
 - 256MB battery backed up cache
 - Dual SCSI320 channels for performance
 - Benchmarks show RAID-5 gives best performance



Other I/O Device Support

- **LAN Attachment**
 - **Multi-ported Ethernet adapters: 10/100/1000**
 - Use caution with Gigabit adapters
- **Tape Drives**
 - **3480, 3490(E), 4mm DAT, SDLT, etc.**
 - **SCSI, Parallel channel, ESCON channel**
- **Channels (Parallel & ESCON)**
 - **Any channel attached devices**
 - printers, terminal ctlrs, 37x5, tapes, etc.
- **ICA (Integrated Communications Adptr)**
 - **6 BSC/SDLC lines as ICA (RS232)**
- **Network attached printers**



Extended Capabilities

- **Shared DASD**
 - Between instances on one system
 - Between two different zFrame systems
 - across TCP/IP connection
- **CTC Connections**
 - Between instances on one system
 - has been used for "basic sysplex" implementation
 - Between instances on two different systems
- **FakeTape**
 - Emulate tape on disk files
 - Via NFS to another system
- **Network Channels**
 - Access devices (i.e. tape) on one FLEX system from another FLEX system via IP connection



FLEX-ES - zFrame



- The zFrame uses *FLEX-ES* as part of a total integrated solution for 'small' S/390 users.
- Other components include
 - IBM xSeries eServers
 - Unix or Linux in customized configuraiton
 - Other I/O (tapes, adapters, etc.)
 - Build & configuration services
 - On-Site Installation & Training
 - On-Going Support



Objective: Provide a product, services, and support, equivalent to a new IBM processor.



Limitations

- **There are some limitations**
 - FICON channels - nothing announced
 - Parallel Sysplex (no CF or timer)
 - No Hardware encryption support
 - 64 bit zSeries support is currently only for PWD developers
 - No QDIO emulation / support



Recent Enhancements

- **Faketape library (AFLIB)**
 - zOS: integrated into SMS and RMM
 - Multiple tape libraries from multiple systems
 - zVM and VSE support under development
- **Compressed faketape**
 - faketape optionally compresses before writing to file
 - Explicit control, or turn on/off as IDRC
 - Supporting utilities

Consider writing compressed faketape files to an automated tape library located on an NFS server.



Recent Enhancements (2)

■ Control Unit Behavior

- ▶ Implemented in the CSI **zCenter**
- ▶ Utilizes **FLEX-ES I/O capabilities**
- ▶ Provide Disk/Tape/Network devices to any "mainframe"
 - Via ESCON channels to traditional zSeries
 - via Network Channels to other FLEX/zFrame systems
 - Use AFLIB to implement a Virtual tape system
 - Emulate disks using xSeries hardware
 - Many 'interesting' possibilities



Futures.....

■ Where do we go from here....?

- ▶ Faster Pentium Processors
- ▶ 64 bit (IA64, EM64T, or ??) processors
 - "64 on 64"
- ▶ Changes to IBM Software Licensing
 - Allow 64 bit mode for Commercial users
 - Consistant Software licensing methodology
- ▶ Continued FLEX-ES enhancements
 - SYSPLEX support(?)
 - ?



Additional References (1)

■ IBM Redbooks

- ▶ **SG24-6215:** "NUMA-Q Enabled for S/390: Technical Introduction"
 - Still the best overall introduction to FLEX-ES
 - Skip over NUMA-Q (xSeries 430) specific sections
- ▶ **SG24-6501:** "S/390 PWD Netfinity enabled for S/390"
- ▶ **SG24-6507:** "S/390 PWD ThinkPad Enabled for S/390"
- ▶ **SG24-6834:** "S/390 Partners in Development: EFS Systems on a Linux Base"
- ▶ **SG24-7007:** "EFS Systems on a Linux Base: Getting Started"
- ▶ **SG24-7008:** "EFS Systems on a Linux Base: Additional Topics"
 - Excellent for users who want to get a little deeper
 - Not just for Linux based systems



Additional References (2)

■ Whitepaper

- ▶ "Exploring the I/O Performance Characteristics of Intel Based FLEX-ES Servers for z/OS" by Dr. H. Pat Artis, Performance Associates, Inc.
http://www.perfassoc.com/flex-es_io_performance_02.pdf
 - This was on an old server; recent tests on new model shows approx 2x this performance.

■ Support Listserve (FLEX-ES)

- www.listserv.uga.edu to subscribe

■ Web pages

- www.csihome.com (and take the zFrame link)
- www.funsoft.com (Fundamental Software Inc)