

zFrame: Exploiting FLEX-ES for small mainframe users.

zFrame/FLEX-ES: Technical Overview and *Update*

NaSPA: NaSTEC²⁰ Conference

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 150 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.

NaSTEC Conference, Oct. 2006



FLEX-ES Architectural Compliance

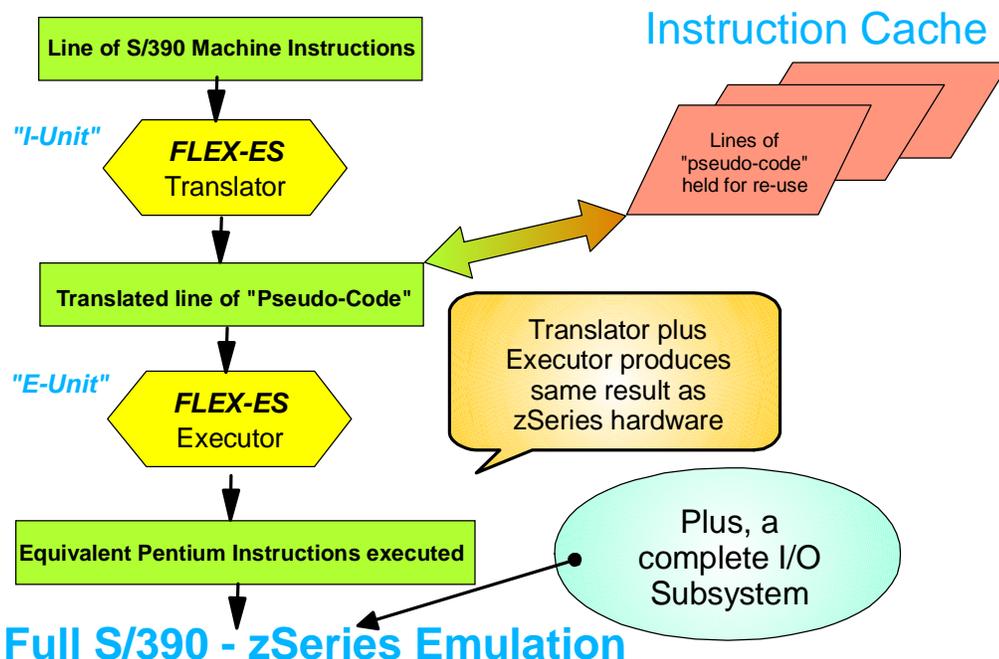
- Meets all IBM definitions for the zSeries architecture (through ALS3)
 - Subjected to (and passed) the same verification tests as new IBM hardware
- Fully compatible with all current IBM operating systems and products
- Recommended by IBM for developers of new vendor software products

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

NaSTEC Conference, Oct. 2006



How does it work?



NaSTEC Conference, Oct. 2006



IBM CMOS - FLEX-ES similarities

■ IBM Systems

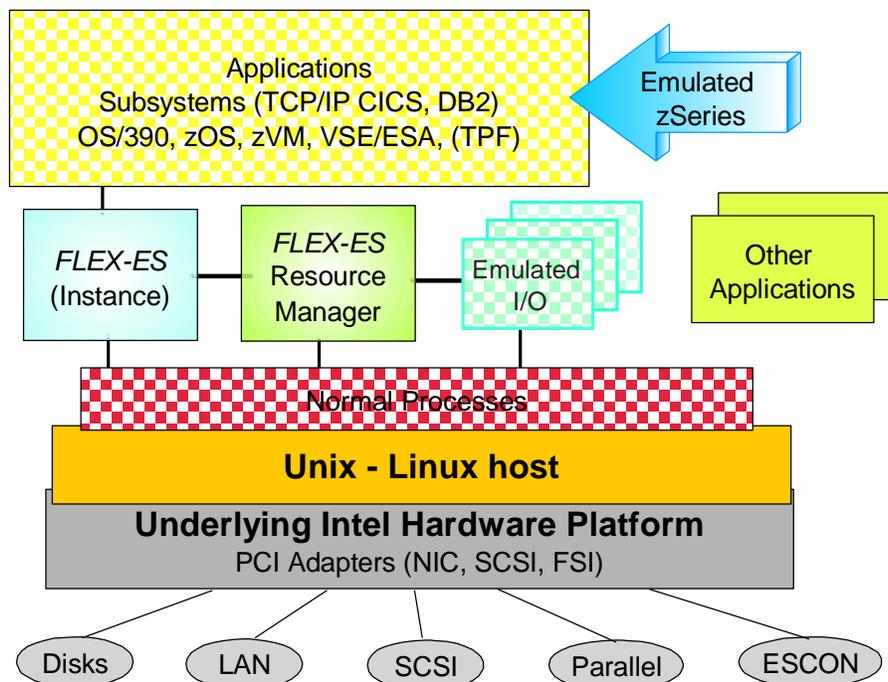
■ FLEX-ES

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

NaSTEC Conference, Oct. 2006



System Structure



NaSTEC Conference, Oct. 2006



Processor Performance

■ Current Processor Performance

- ▶ **Intel Pentium technology allows (approx):**
 - Single processor systems up to 37 MIPS
 - Dual processor systems up to 72 MIPS
 - Three processor systems up to 110 MIPS
- ▶ **AMD Opteron**
 - Single processor up to 68 MIPS
 - Dual/Triple processor of 130 and 180 MIPS
 - Four Dual-Core chips could reach 400 MIPS
- ▶ **Intel "Core-Duo" not yet benchmarked**
 - Expect performance similar to Opteron

NaSTEC Conference, Oct. 2006



Processor Architecture

■ S/370, ESA, or 64 bit zSeries mode

- 64 bit currently restricted to PWD systems
 - Runs z/OS 1.7 and zVM 5.2
 - Commercial systems currently limited to zOS 1.5, zVM 4.4, and VSE 3.1

■ Up to 11 GB of memory for S/390

- Maximum of about 1.8 GB per "instance" for FLEX-ES ver 6 or 7
- FLEX-ES Ver 8 supports > 2 GB memory size (for 64 bit emulated systems)
- Larger memory sizes (>11 GB) are possible, but not generally practical

NaSTEC Conference, Oct. 2006



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 | 30 | 4.0 GB | 330 GB | 2/1 | 236 2 x 3.2 GHz |
| z60 | 60 | 4.0 GB | 330 GB | 2/2 | 236 2 x 3.2 GHz |
| z90 | 110 | 6.0 GB | 330 GB | 4/3 | x3800 4 x 3.6 GHz |

NaSTEC Conference, Oct. 2006



IBM Software License Charges

| MIPS | License Base | Comments |
|-------------|---------------------|---|
| 8 | ESL | ESL withdrawn by IBM: may be available for current ESL users |
| 16 | 1/3 GOLC | Attractive pricing. Limited to Pentium III 1.0 GHz |
| 30 | 2/3 GOLC | Fastest commercial Uni-Processor Less cost than 36 MIPS z800 |
| 60 | GOLC - H30 | Same license as MP3000 H30 |
| 110 | GOLC - H50 | Same license as MP3000 H50 |

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

NaSTEC Conference, Oct. 2006



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

NaSTEC Conference, Oct. 2006



Current zDev Modelsh

■ Standard zDev base configurations

| | MIPS (approx) | S/390 Memory | Disk (GB) | # of Proc. | xSeries |
|--------------|---------------------|-----------------|--------------|---------------|-------------------------------|
| zPad | 40 (I/O Limited) | 1.5 GB | 80 GB | 1/1 | TP T60 |
| zDev1 | 36 | 4 GB | 160 GB | 2/1 | 236 2 x 3.6 GHz |
| zDev2 | 65 | 4 GB | 330 GB | 2/2 | 236 2 x 3.6 GHz |
| zDev3 | 110 | 6.0 GB | 330 GB | 4/3 | 260 4 x 3.6 GHz |
| zDevY | 130 | 6.0 GB | 420 GB | 2 / 2 | Intellistation 2 x 2.8 GHz |

NaSTEC Conference, Oct. 2006



More on Opteron zDevY

- **Not available in IBM xSeries "Industrial Strength" Servers**
 - Not suitable for Commercial / Production
 - May be appropriate for development use
 - "Server Quality" should be available soon
- **Single and dual Core processors up to 2.8 GHz**
 - 2.8 GHz dual core has not been benchmarked
 - About 68 MIPS per processor (single core)
- **Using high performance SCSI RAID disks for up to 840 GB internal disk**

NaSTEC Conference, Oct. 2006



zFrame 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
 - Over 6 TB very easy and practical
 - Very economical cost point
- **Disk Performance**
 - Very good performance: 1K - 3K IO / sec
 - 20 - 30 *microseconds* for cache read hit
 - See http://www.perfassoc.com/jsc/pdf/papers/flex-es_io_performance_02.pdf
- **RAID-5 Disk array (IBM ServeRAID adapters)**
 - Full data redundancy
 - 256MB battery backed up cache
 - Dual SCSI320 channels for performance
 - Benchmarks show RAID-5 gives best performance

NaSTEC Conference, Oct. 2006



Recent Enhancements

- **Faketape library (AFLIB)**
 - And related new functions
- **Compressed / Encrypted faketape**
 - Opens up new possibilities
- **New SuSE Linux as base**
 - Improved flexibility and support
- **FLEX-CUB & the zCenter**
 - A complete new way to exploit FLEX technology
- **FLEX-ES Ver 8**
- **Remote tape Vault (zVault)**

Let's look at these in more detail.....

NaSTEC Conference, Oct. 2006



FakeTape Library: AFLIB

- **zOS: integrated into SMS and RMM**
 - High function, Tape management system
 - VTS (Virtual Tape System) type capability
 - Tested on and packaged for zOS 1.5 and 1.6
- **OS390: Full Integration not available**
 - Basic 'automounter' mode only
 - Tested on and packaged for OS/390 2.10
- **Features**
 - Multiple tape libraries from multiple Instances
 - "Tape library" can be on same system or across channel or NFS connection
 - Library can be on a FLEX-CUB control unit
- **Requirements:**
 - FLEX-ES Ver, 7(+)
 - Extra cost option

NaSTEC Conference, Oct. 2006



fsihost for zVM

■ Allows zVM users to issue commands to the host Linux/Unix system

- can send commands to linux to manipulate files
 - rename (mv) , remove, etc.
 - `fsihost (cmd ls -l`
 - `pipe cms fsihost dev 7ff (cmd cat myfile | > my file a V`
- Can send commands via linux to FLEX-ES command line processor
 - `fsihost (cmd echo mount 580 /faketape/new-file | flexescli localhost zvm`
- Allows automation and control of faketape mounting (or other activities) from within zVM
- Use 'enabling EXECs' to simplify further
 - `flex mount 580 /faketape/other-file`

■ Requirements:

- Free download for FLEX-ES and FLEX-CUB users

NaSTEC Conference, Oct. 2006



Compressed FakeTape

■ FakeTape files can be automatically compressed as they are created

- Multiple controls on which tape files are compressed
- Compression is done at the Flex host level
 - Pentium cycles rather than zSeries cycles
- Transparent to zSeries operating system
- Consider:
 - writing compressed tape files across an NFS connection to a remote server
 - Combining Compressed FakeTapes with AFLIB

■ Requirements:

- FLEX-ES Ver, 7(+)

NaSTEC Conference, Oct. 2006



Encrypted FakeTape

- **Important topic in today's IT industry**
- **Provided by emulated tape device in FLEX 7.0.7 and later**
 - **Selectable OpenSSL AES128, 192, 256**
 - **Specifiable (encrypted) key files**
 - **Can write encrypted files to local disk or FSI's new remote DR facility**
 - **FLEX tape related utilities updated to support encryption**
 - **AFLIB controlled tapes can be encrypted**
- **FAKETAPE encryption only**
 - **not for SCSI or channel attached tapes**

NaSTEC Conference, Oct. 2006



SuSE Linux as base

- **Moving to SuSE Linux**
 - **Currently 9.3 Pro: soon 10.x**
 - **SuSE, SLES, SLED**
 - **Better support than previous RedHat**
 - **More current version of drivers and kernel**
 - **Positions us better for move to 64 bit**
 - **UnixWare 7.1.4 continues to be available**

NaSTEC Conference, Oct. 2006



CUB: Control Unit Behavior

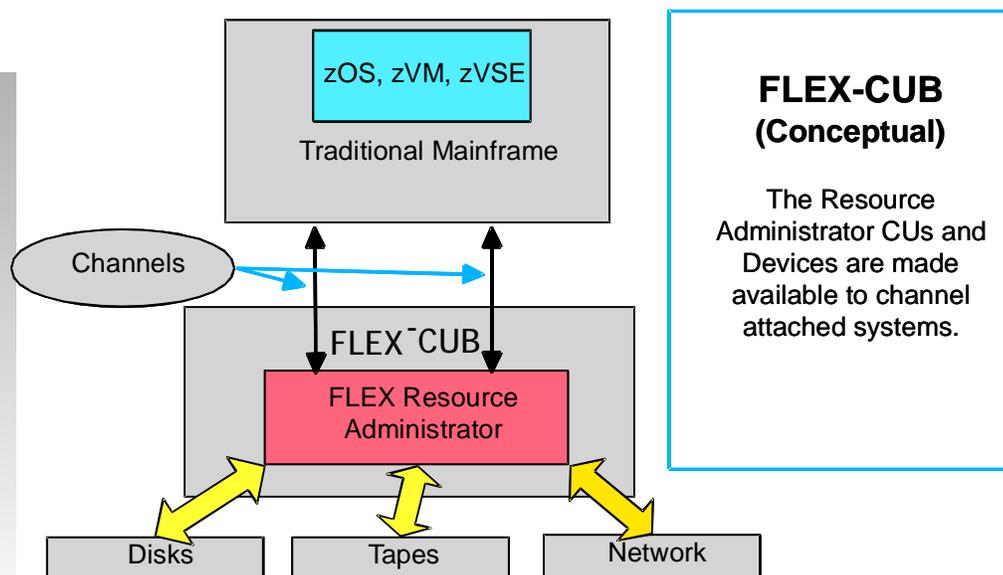
■ Implemented in the CSI zCenter

- ▶ Utilizes FLEX-ES I/O capabilities
- ▶ Provide Disk/Tape/Network devices to any "mainframe"
 - Via ESCON/Parallel 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

NaSTEC Conference, Oct. 2006



Conceptual View of CUB



NaSTEC Conference, Oct. 2006



Storage Server

- **Emulate multiple types of disk devices**
 - ▶ All are RAID-5 protected using IBM high performance disks and RAID controllers
- **Capacities from 330 GB to over 6 TB**
- **Transparent, tunable, and very effective disk cache (1 – 3 GB of memory)**
- **One or multiple channels to connect to one or more clients (DASD can be shared)**
- **Can be configured as multiple Logical CUs**

NaSTEC Conference, Oct. 2006



Tape Server

- **Two Major kinds of tape service**
 - ▶ **SCSI tape server** allows connection of SCSI tape devices to mainframes as traditional 34x0 devices
 - Either media compatible 3480/90 or new technology such as SDLT (320+ GB in one Cartridge)
 - ▶ **FakeTape server** allows creation of a Virtual Tape System with up to 6 TB of online storage
 - FSI provided software to support tape library functions in zOS (AFLIB) and zVM. (OEM support for VSE)
 - Backup/offload of tape files via included SDLT tape drive
 - Tape files can be compressed for even greater capacity

NaSTEC Conference, Oct. 2006



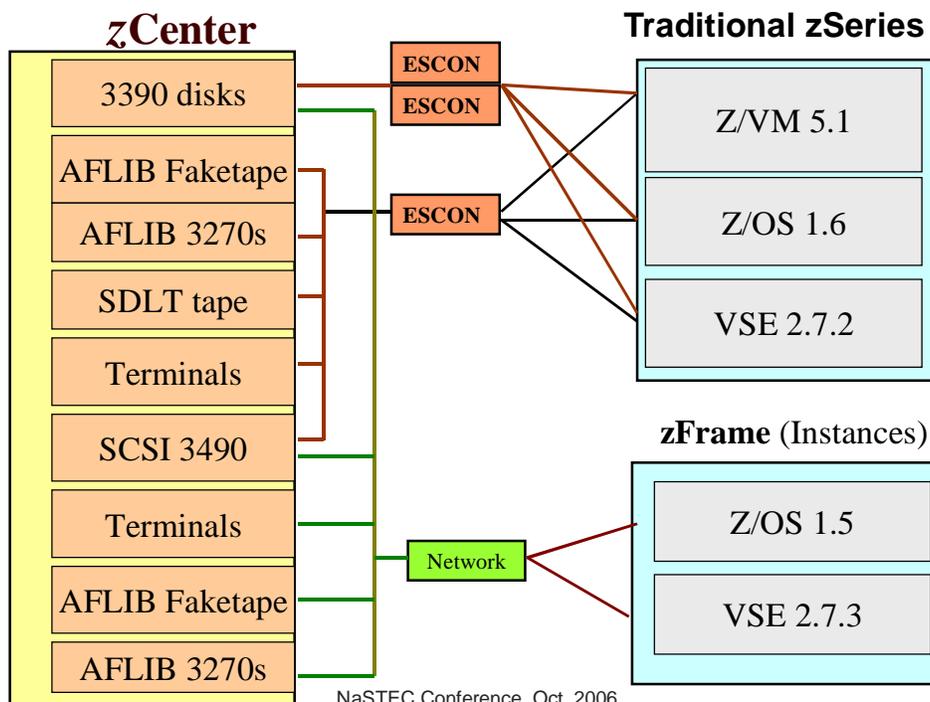
Terminal – Network Server

- Connect your IP attached workstations to your zSeries without using TCP/IP stack cycles
- Makes IP connected workstations appear as local, non-SNA 3270 devices on a 3174
- Can be “EMIF’ed”, providing operator console support to multiple LPARs (as IBM 2074) replacing multiple 3174s.
- IP attached workstations utilize standard “TN3270” type terminal emulators

NaSTEC Conference, Oct. 2006



Complex zCenter Example



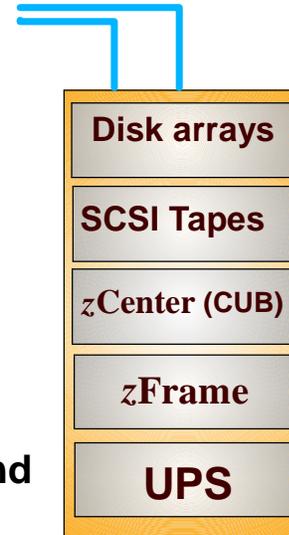
NaSTEC Conference, Oct. 2006



Full zCenter Configuration

- **Combine FLEX-CUB and standard zFrame in one package**
 - Provide I/O services to channel attached traditional mainframe
 - With full, normal zFrame processor capabilities
 - In one compact rack mounted package
 - Share DASD between zFrame and zSeries

Channels to zSeries



NaSTEC Conference, Oct. 2006



FLEX-ES Ver 8

- **"64 on 64": 64 bit zArchitecture emulation on 64 bit AMD/Intel systems**
 - Utilizes Intel EM64T or AMD 64 support
 - Allows large memory sizes
 - > 2GB for 64 bit instances
 - Full 2 GB for 31 bit instances
 - Improved processor performance
 - Some improvement for 31 bit mode
 - Big improvement for 64 bit mode instructions
 - Complete re-write of processor emulation
- **Status:**
 - Now available
 - Subject to existing licensing restrictions

NaSTEC Conference, Oct. 2006



zVault

- Uses **FLEX-ES** remote tape vault abilities
- Backup to faketape (or zCenter) goes directly to FS Vault location
- Transparent to zSeries system
- Speed dependent on Internet Connection
- Data is encrypted during transmission and at vault site
- Automatic data compression



NaSTEC Conference, Oct. 2006



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 configuration
 - 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 (none)**
 - ▶ **Parallel Sysplex (no CF or timer)**
 - ▶ **No Hardware encryption/compression**
 - ▶ **64 bit zSeries support is currently only for PWD developers**
 - ▶ **No QDIO emulation / support**

NaSTEC Conference, Oct. 2006



Futures.....

- **Where do we go from here....?**
 - ▶ **Faster & Dual Core Intel Processors**
 - ▶ **Exploit AMD processor family**
 - Commercial/production quality servers
 - ▶ **64 bit (IA64, EM64T, or AMD) processors**
 - "64 on 64" (available in FLEX-ES Ver 8)
 - ▶ **Changes to IBM Software Licensing**
 - Allow 64 bit mode for Commercial users
 - Consistant Software licensing methodology
 - ▶ **Continued FLEX-ES enhancements**
 - z9 instructions
 - ?

NaSTEC Conference, Oct. 2006



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

NaSTEC Conference, Oct. 2006



Additional References (2)

■ Whitepapers

- ▶ "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.
- ▶ "The Cornerstone zFrame: An Overview" A general overview of FLEX-ES and the zFrame
 - Cornerstone Web pages or our exhibit stand

■ 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)

NaSTEC Conference, Oct. 2006



The End...

Thank you!

NaSTEC Conference, Oct. 2006