



Oracle 10G– System z Linux

May 2008 Oracle Special Interest
Group

Boston University
Gerard C. Shockley

Luke Mcgee

gshock@bu.edu

lcmgee@bu.edu

www.bu.edu



Why Oracle zSeries Linux at BU:

- ❑ Where we started – Academic Curiosity – CPU2000 Specint benchmarks 9672-G3 IBM (and others) 2002 timeframe things actually worked. Emulated FP issues
- ❑ Started with SLES 7 – Bare Metal ; then zVM 4
- ❑ Got involved in the IBM Linux Community Development
- ❑ Executed a Successful platform consolidation from Windows Oracle to DB2 on System z Linux (PROD 2004)
- ❑ Excellent track record with vended zLinux offerings (IBM,RSD)
- ❑ Open Source contributions (To name a few tools)
 - *TOMCAT, ANT, JAVA, POSTGRESQL, PAM,*
 - *PERL, PYTHON, JAVA, GCC, STRACE, RCS, WGET, MYSQL, PHP*
- ❑ Leadership support for virtualization (vFirst Philosophy)
- ❑ Tight integration with the platform (Hipersockets)
- ❑ IBM contributions and leadership in Linux on z (SLES)
- ❑ Built 3 tier configurations



History of z/Linux

- ❑ 1998 Marist College download enabled
- ❑ 1999 IBM releases some OCO to OSS community
- ❑ 2000 SuSE , Turbolinux Announced
- ❑ 2001 Redhat announced
- ❑ 2002 64bit distributions announced Debian
- ❑ 2004 IBM Build your own – ibm.com/developerworks
 - Kernel 2.6.5 s390-8 currently available
 - Binutils 2.15
 - gcc 3.4.0
 - glibc 2.3.2
 - gdb 6.1, strace 4.5.2,
- ❑ 2004 BU goes live EOS.Thinclient
- ❑ 2006 First uPortal.org my.bu.edu to run on zLinux
- ❑ 2007 BU goes live Oracle 10G on zLinux



Why z/Series

- ❑ z/Series hardware – 99.999% uptime z means zero downtime
Industrial Strength Hipersockets technology, MTBF Measured in decades, Dual instruction execution, Redundancy, ECC
- ❑ z/VM – Virtualization Technology
- ❑ Hardware Cryptographic Support – Accelerator for TLS,SSL
- ❑ Infrastructure Simplification – Server Consolidation of lightly used servers into a single z/Series complex
- ❑ Potential for TCO Savings -
Hardware/Software/Middleware/Systems Management/Staff/
- ❑ Bi-modal 31 and 64 bit distributions
- ❑ IFL – Engine support – Dedicated to LINUX
- ❑ HiperSockets



z/Linux: The Operating System

- ❑ Kernel: Low level system control and interfaces, program and hardware device drivers
- ❑ Init: Main program executed post kernel load. Supports multiple run levels
- ❑ File Structure: EXT2, EXT3, RieserFS, JFS, NFS, SWAP, PROCFS, SMBFS
- ❑ Applications: Many categories, ex: OpenOffice, web, data and systems management , scientific, portable, government systems, transactional

- ❑ Open Source contributions (To name a few tools)
 - PERL, PYTHON, JAVA, STRACE, RCS, WGET, MYSQL
 - LIBTOOL, LIBCA, LIBTIFF, KDE, CUPS, SASH,
 - IPTABLES, DHCP -TOOLS, BZIP, IMAP, s/390 Crypto tools,
 - TCL, ZSH, GDK, TCPDUMP, IPRROUTE, SYSSTAT, JSERV, MOD-PERL, LDAP, OPENSSSH, BINUTILS, SAMBA, UCDSNMPD
 - OPENSSSL, KERNEL SOURCE, RPM, MOZILLA

Systems Tasks Leading up to Production Activation



- ❑ Think/Design/Build/Clone
- ❑ Infrastructure Design – Guest Creation 3 tier
- ❑ Guest directory – CPUs/Memory/Minidisks
- ❑ Cloning considerations /usr /opt /var
- ❑ Network/Security Infrastructure design – VCTC/Guest LAN QDIO/Vswitch – Good redbooks as training source.
- ❑ File System Structures – Determine types and mount points
- ❑ Operating System installation and formatting
- ❑ Application installation and tailoring
- ❑ Disaster Recovery planning and validation
- ❑ Performance and Capacity planning – z/VM and z/Linux
- ❑ Tooling exploitation – Mix of OSS/vended components
- ❑ Debugging – z/Series specific (strace, SysRq, top, s390dbf)
 - Kernel debugging – RAS, LTT – tracing toolkit, dprobes, kdb
 - Event logging – standard logs and vmlogs guest logs

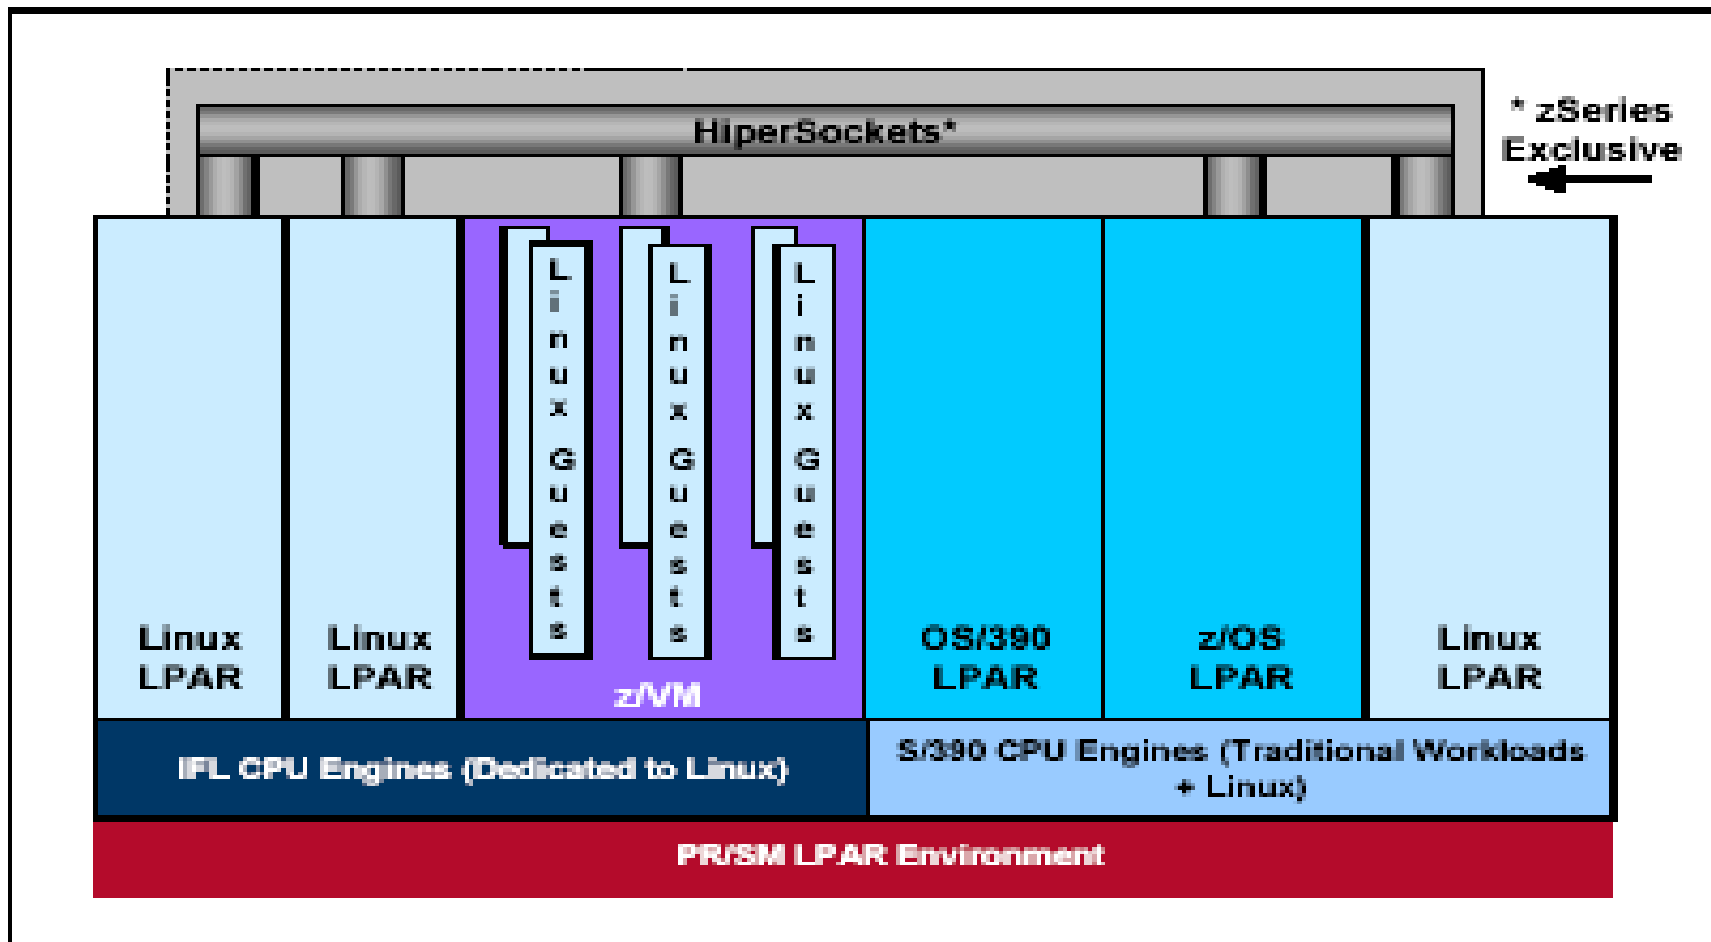


Primary Business Challenges

- ❑ Economic Factors
 - Entry Costs associated with Vended Software, ongoing costs
 - Costing estimates remaining within project scope
 - Mix of traditional Vended and Community based Sources (OSS)
- ❑ Reduce project cycle times versus d-deploy options (Clones)
- ❑ Application integration with existing mixed technologies zOS
- ❑ Minimizing staff rebuild effort (clone approach)
 - Vended Software – Exit points , OSS – Source modifications
- ❑ Staff management overhead (build once use many)
 - Care and feeding, versus the vended support model
- ❑ Ability to support with current or fewer teams
- ❑ Centralizing server management (zSeries)
- ❑ Maintain current security Models and strength
- ❑ Lowering power consumption
- ❑ Simplification of Disaster Recovery



The System z Complex





BU Projects System z Linux

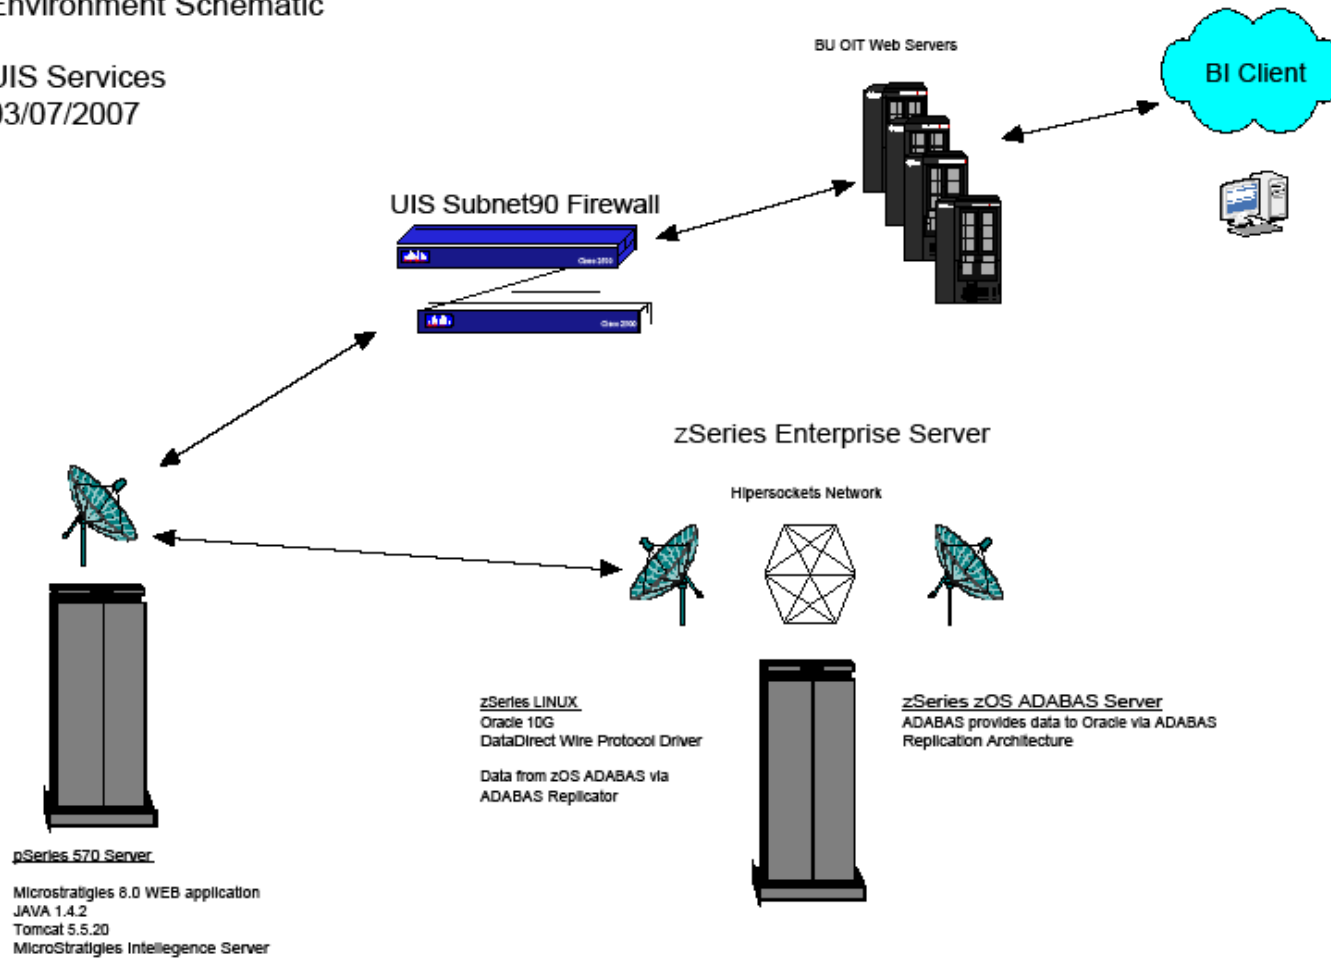
- ❑ Oracle 10G Projects Prod 10/2007
 - Business Intelligence BU-DAR (PROD 10/2007)
 - ❑ Data WHSE built using OWB and database replication with Java – XML utilities.
 - ❑ Datadirect thin wire protocol driver – AIX – zLinux
 - ❑ RMAN operational configuration with FDR full , upstream incremental – looking flash
 - University Document Imaging Project (Scanning, retrieval, workflow)
 - ❑ Onbase System Selected
 - ❑ Target Oracle 10G
 - ❑ Enterprise Wide System
 - ❑ Stress showed 18% CPU for a single guest
 - ❑ Integrated with zOS systems
 - Server Consolidation (Proof of concept)– Windows to zSeries
 - ❑ Coeus – Open source Grant and Contract Management (POC) (OCI)
 - ❑ INSPIR – Medical Software System
 - ❑ Oracle DBA BU Services – BUMC CPU Oracle 9I
 - Enterprise Oracle Grid Server Monitoring Project
 - ❑ I586 Server running OpenSuse OMS with zSeries Linux Agent servers on the backend (Non Prod)
- ❑ IBM HostOnDemand – Java Servlet Emulation Client Prod 08/2004
- ❑ EOS Thin Client – Java based output viewing technology 08/2005
- ❑ Java Enterprise Edition Prod 2005
 - Student Graphics Scheduler – Student Schedule Matrix - *Very successful*
 - BUCHART – Faculty charting tool
 - Schedule Servlet – Student Course schedule matrix
 - ID-Sync Project – Quartz Scheduler, Hibernate JDBC interface
- ❑ Eclipse based development – zSeries Deployment
 - 12F – Java2anywhere




BU Oracle 10g System z Linux

BU Business Intelligence
Environment Schematic


UIS Services
03/07/2007



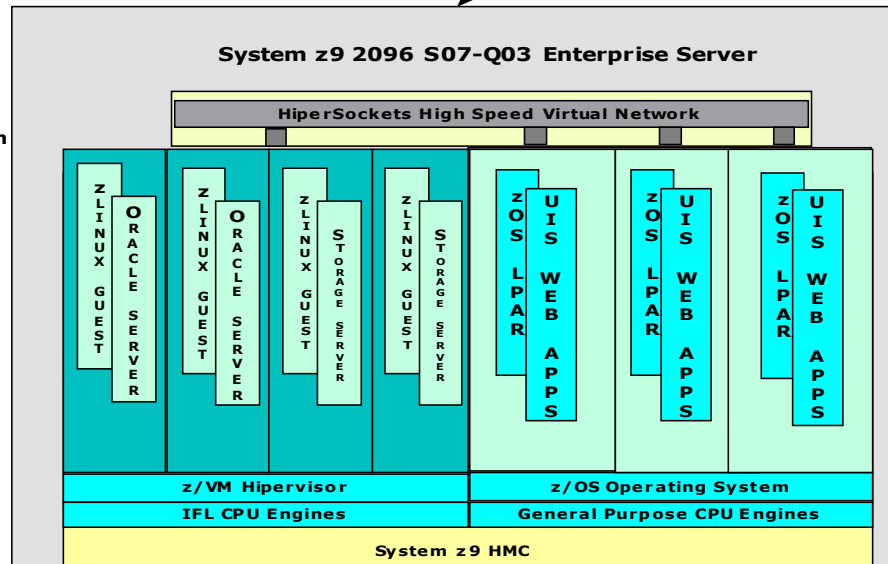
Enterprise Document Imaging



Proposed Document Imaging Configuration Boston University



zSeries SuSE LINUX
 Oracle 10G, DB2, JAVA EE
 High performance communication with zOS via Hipersockets



zSeries zOS ADABAS Server
 ADABAS provide institutional class data store.



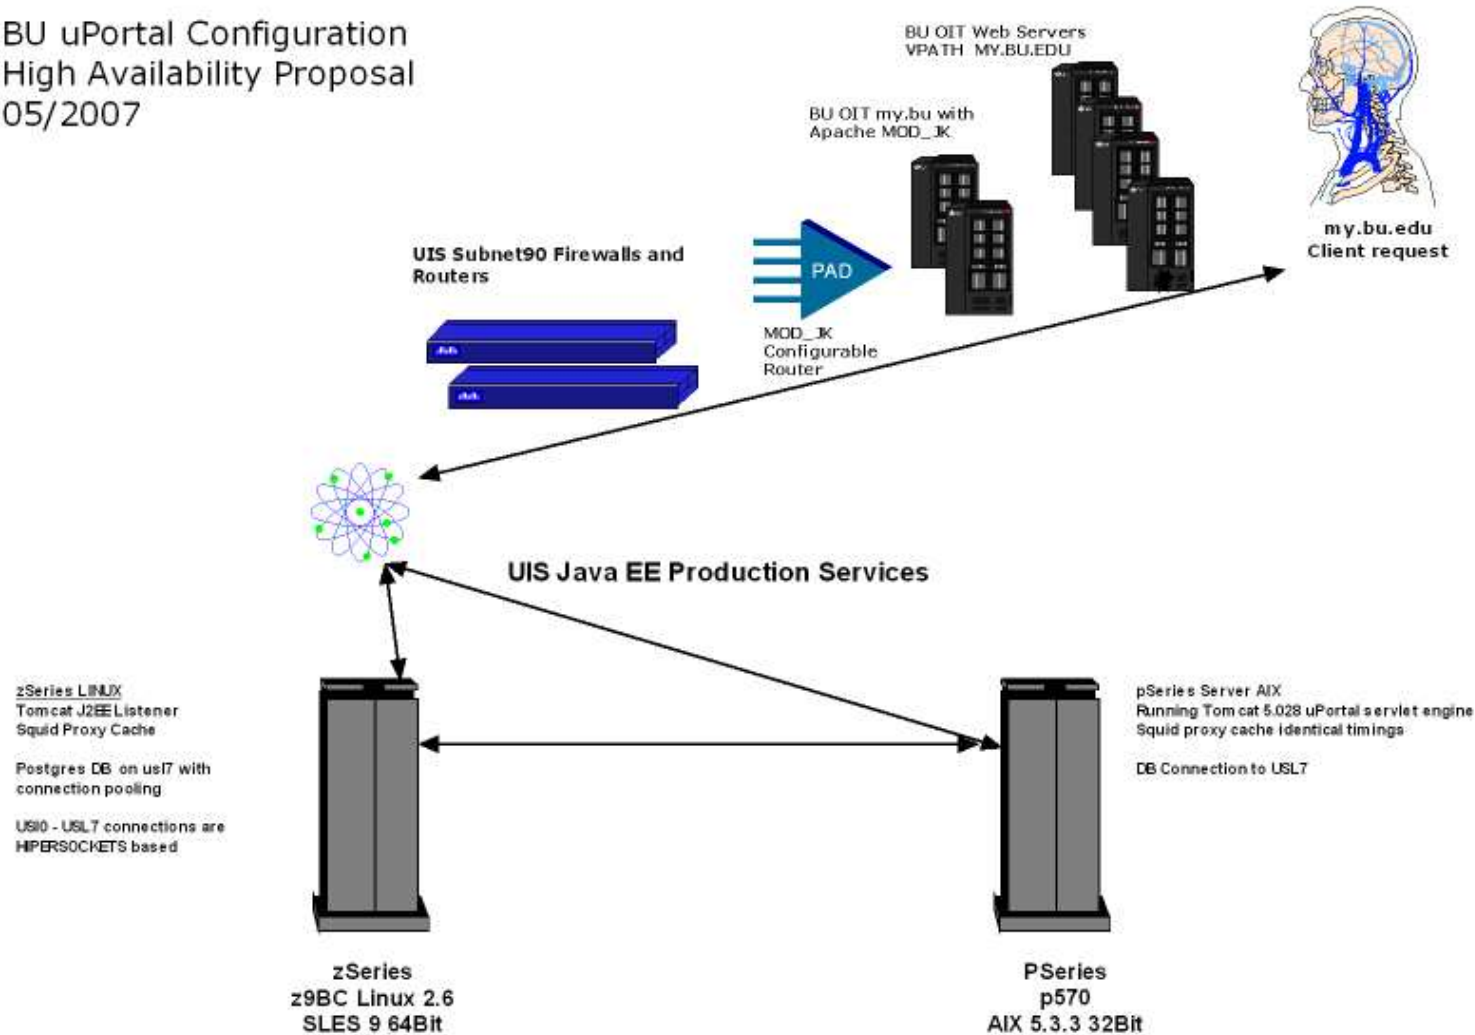
BU z/Linux Open Source projects

- **uPortal** – uPortal.org
 - Java Based OSS Portal
 - Built by JASIG.org Java Open Source Project
 - 1st Sytem z Linux Environment for uPortal
 - Business Drivers
 - Community supported direction – Similar challenges
 - Need for high volume transaction processor for Linux – BlueGeneL used for workload simulations
 - We Installed Tomcat 5.0 – Open Source Application server <http://jakarta.apache.org/tomcat/index.html>
 - Catalina project – Servlet 2.3 and JSP 1.2 specification
 - OSS Database Server environment (PostgreSQL)
 - Success Strategies
 - Involve all appropriate internal organizations – early and frequently
 - Secure appropriate external support organizations – we use Sine Nomine <http://sinenomine.net/> 24x7



BU uPortal Schematic

BU uPortal Configuration
High Availability Proposal
05/2007





Summary zBU

- ❑ Capacity Upgrades without disruption
- ❑ ROCK Solid Virtualization with zVM
 - Simply NEVER a problem in ~ 4 years PROD
- ❑ Excellent Scalable (dynamic) performance – IFL growth, Memory, OSA
- ❑ HiperSockets support – All guests for backups, installs via hipersockets.
- ❑ Server Consolidation Opportunities
- ❑ Cost of Computing Reductions Opportunities

Questions

