#SHAREorg

SHARE
Technology · Connections · Results

Implementation of Red Hat Linux on z: User Experiences at Isracard

Mike Shorkend
mshorkend@isracard.co.il

Tuesday, March 15, 2012
Session Number 10441

Isracard Group

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AIX*
DB2*
HiperSockets
IBM*
IBM logo*
IMS
CICS
System z
System z9
System z10
Tivoli
WebSphere*
z/OS*
z/VM*
zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States, other countries or both.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows, Windows NT and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates. More information on Oracle trademarks can be found at www.oracle.com/html/copyright.html.
Istrobe is a registered trademark of Compuware
CSL-WAVE is a trademark of CSL international
CA-Unicenter and CA Wily Inroscope are trademarks of Computer Associates International
Netvackup is a trade mark odf Symantec.
Control-M is a trademark of BMC

* All other products may be trademarks or registered trademarks of their respective companies.

Agenda

- **Introduction**
- Some History
- Current Status
- Challenges and Solutions
- Observations
- Questions

What you won't hear today

Why Virtualization and Consolidation are good

Linux kernel

Bash

rpm's

LVM

What you *will* hear today

- ✓ Why zLinux was a good choice for us
- ✓ How we are doing it
- ✓ The potholes along the way(and how we fixed them or bypassed them)
- ✓ Which applications were ported
- ✓ Our toolbox
- ✓ Decisions that we might have taken(or not taken) if we had seen this presentation before we started
- ✓ The challenges
- ✓ How to implement zLinux at smaller shops

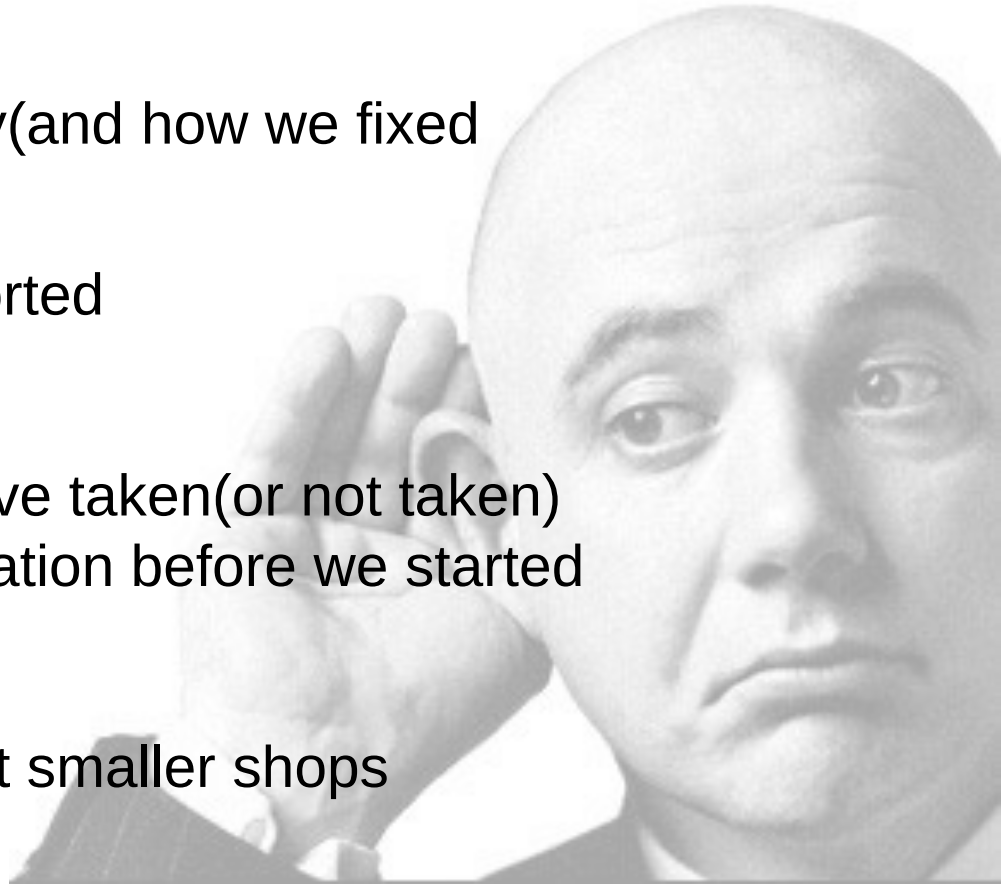

#SHAREorg

SHARE
Technology • Connections • Results

Over 100,000 merchants

Over 50 million business transaction per month

Monthly turnover of 8 billion NIS

3.4 Million Cards
47 % market share

2 million card holders

Isracard Group

Isracard Infrastructure

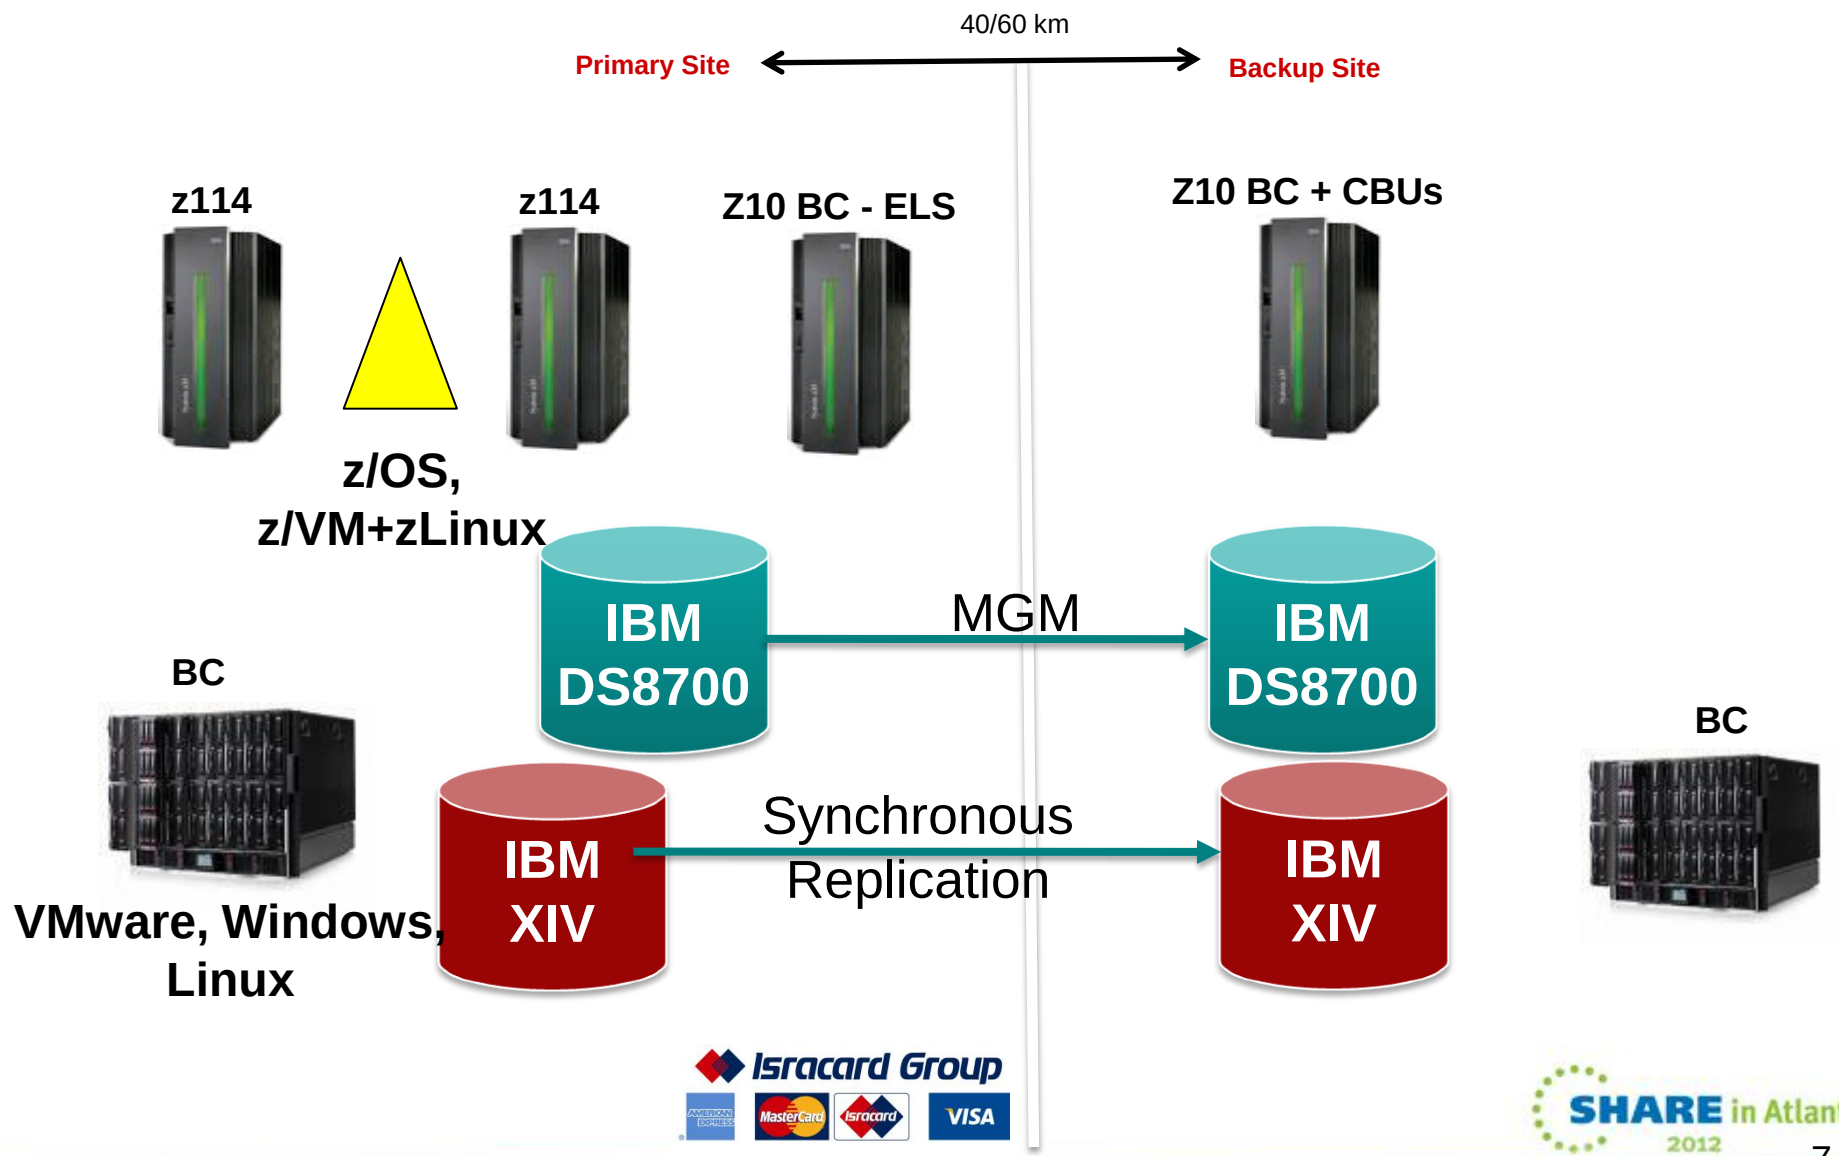

Agenda

- Introduction
- **Some History**
- Current Status
- Challenges and Solutions
- Observations
- Questions

Isracard Before Consolidation

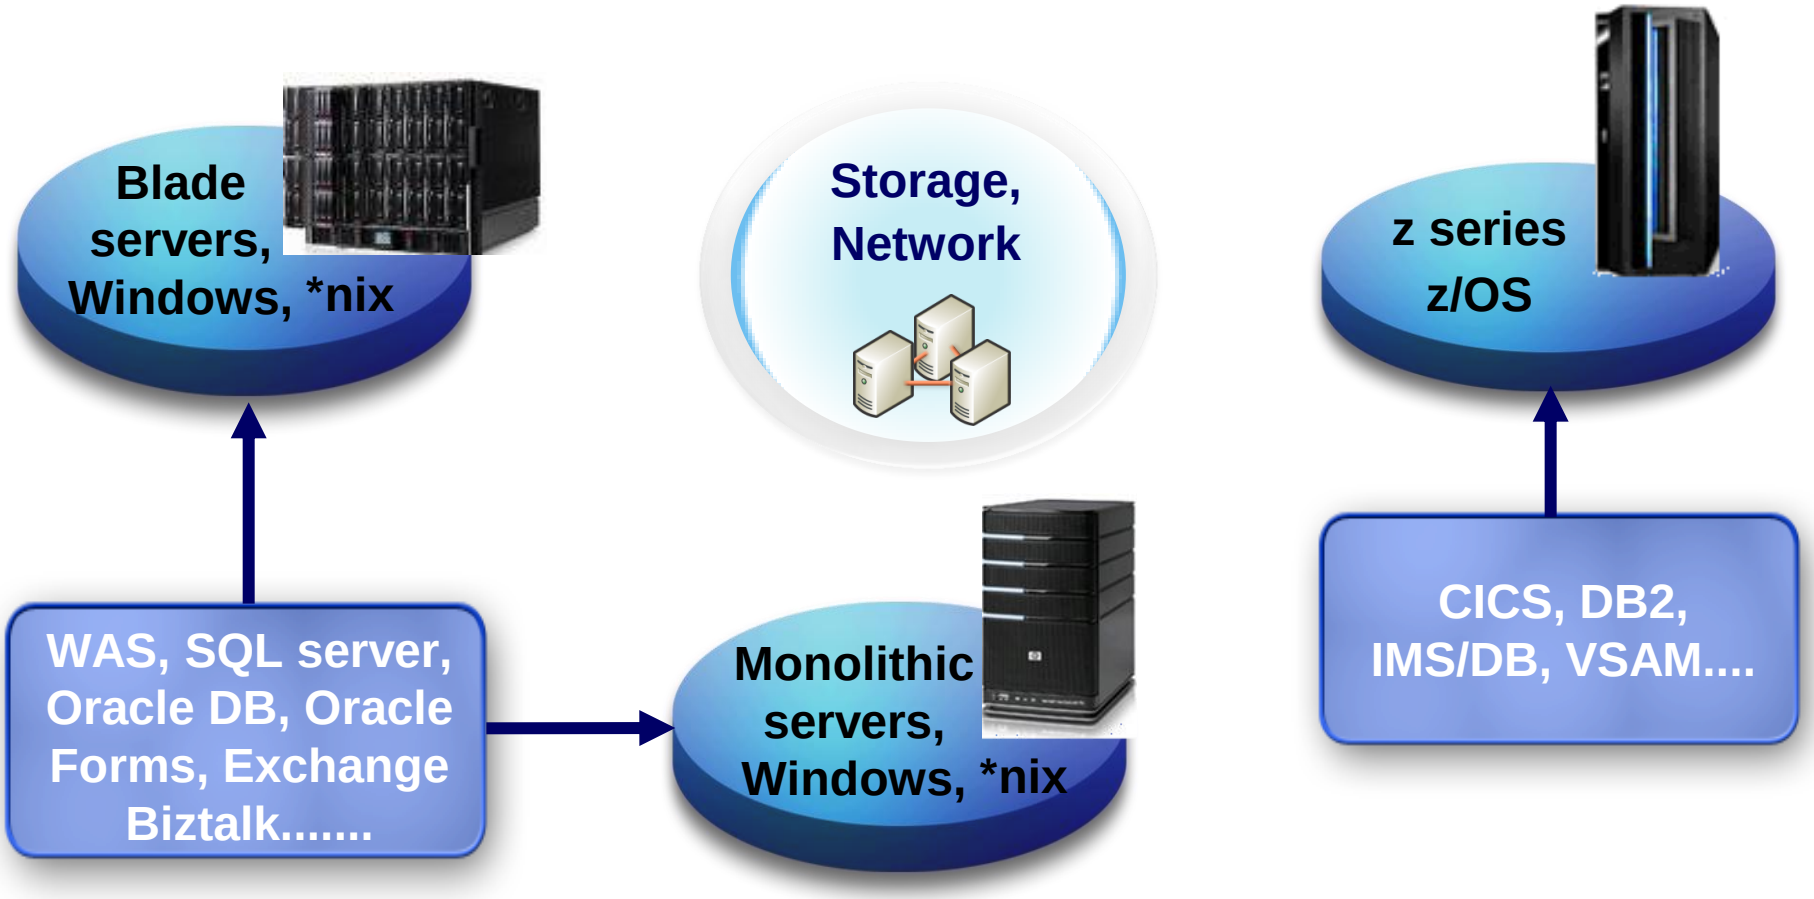

DR Infrastructure before consolidation (3Q08)

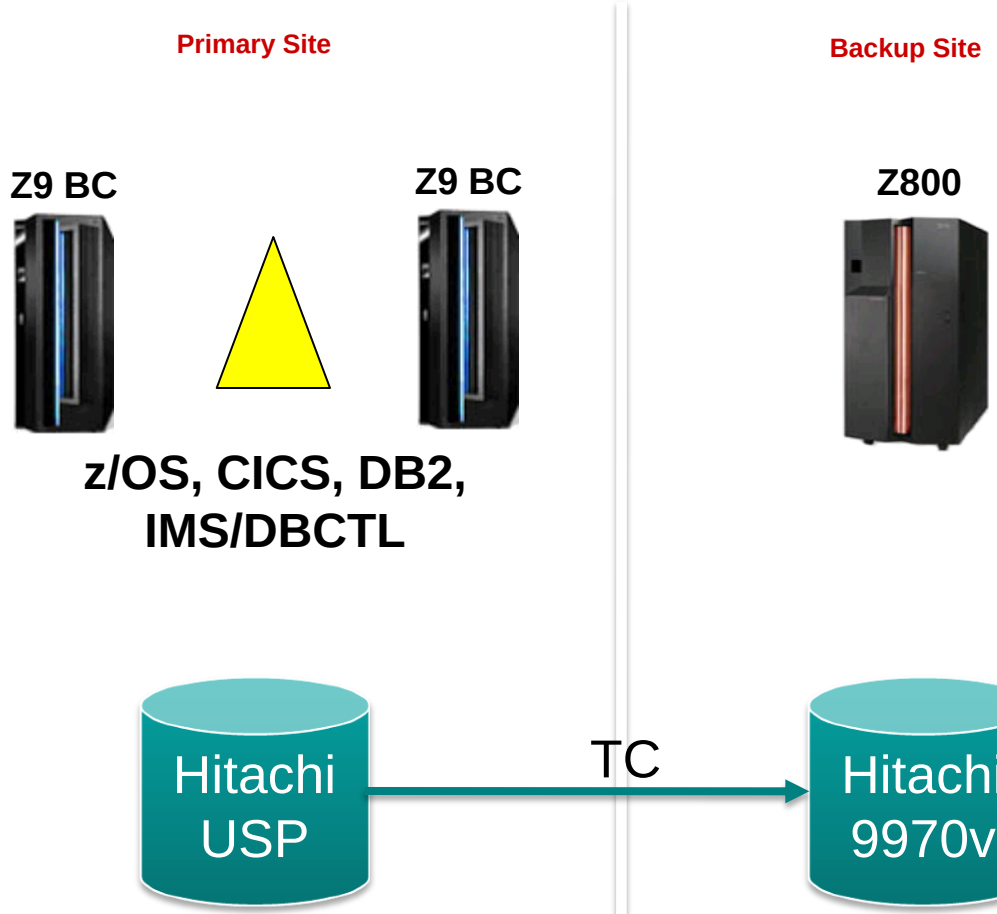

The consolidation trigger

- Until 2008, all core business was on z/OS - hence the distributed systems were not available at the backup site
- Core business on distributed systems - management decision to have them at backup site as well
- Backup site floor space and environmentals were extremely restricted
- We already had a mainframe at the backup site, so zLinux did not take up any additional floor space/power/cooling
- Servers that can not go to zLinux will be consolidated on VMware and blades

Why (z)Linux?

- ✓ Total Cost of Ownership
  - Oracle is the go/nogo
  - We found that the break even point is 1 BC = 1 IFL with 32GB
- ✓ Server Management is easier (see CSL-WAVE later on)
- ✓ Built-in DR
- ✓ RASSS
  - Reliability, Availability, Security, Stability, Scalability
- ✓ Performance
- ✓ Close to the core business

Isracard Infrastructure

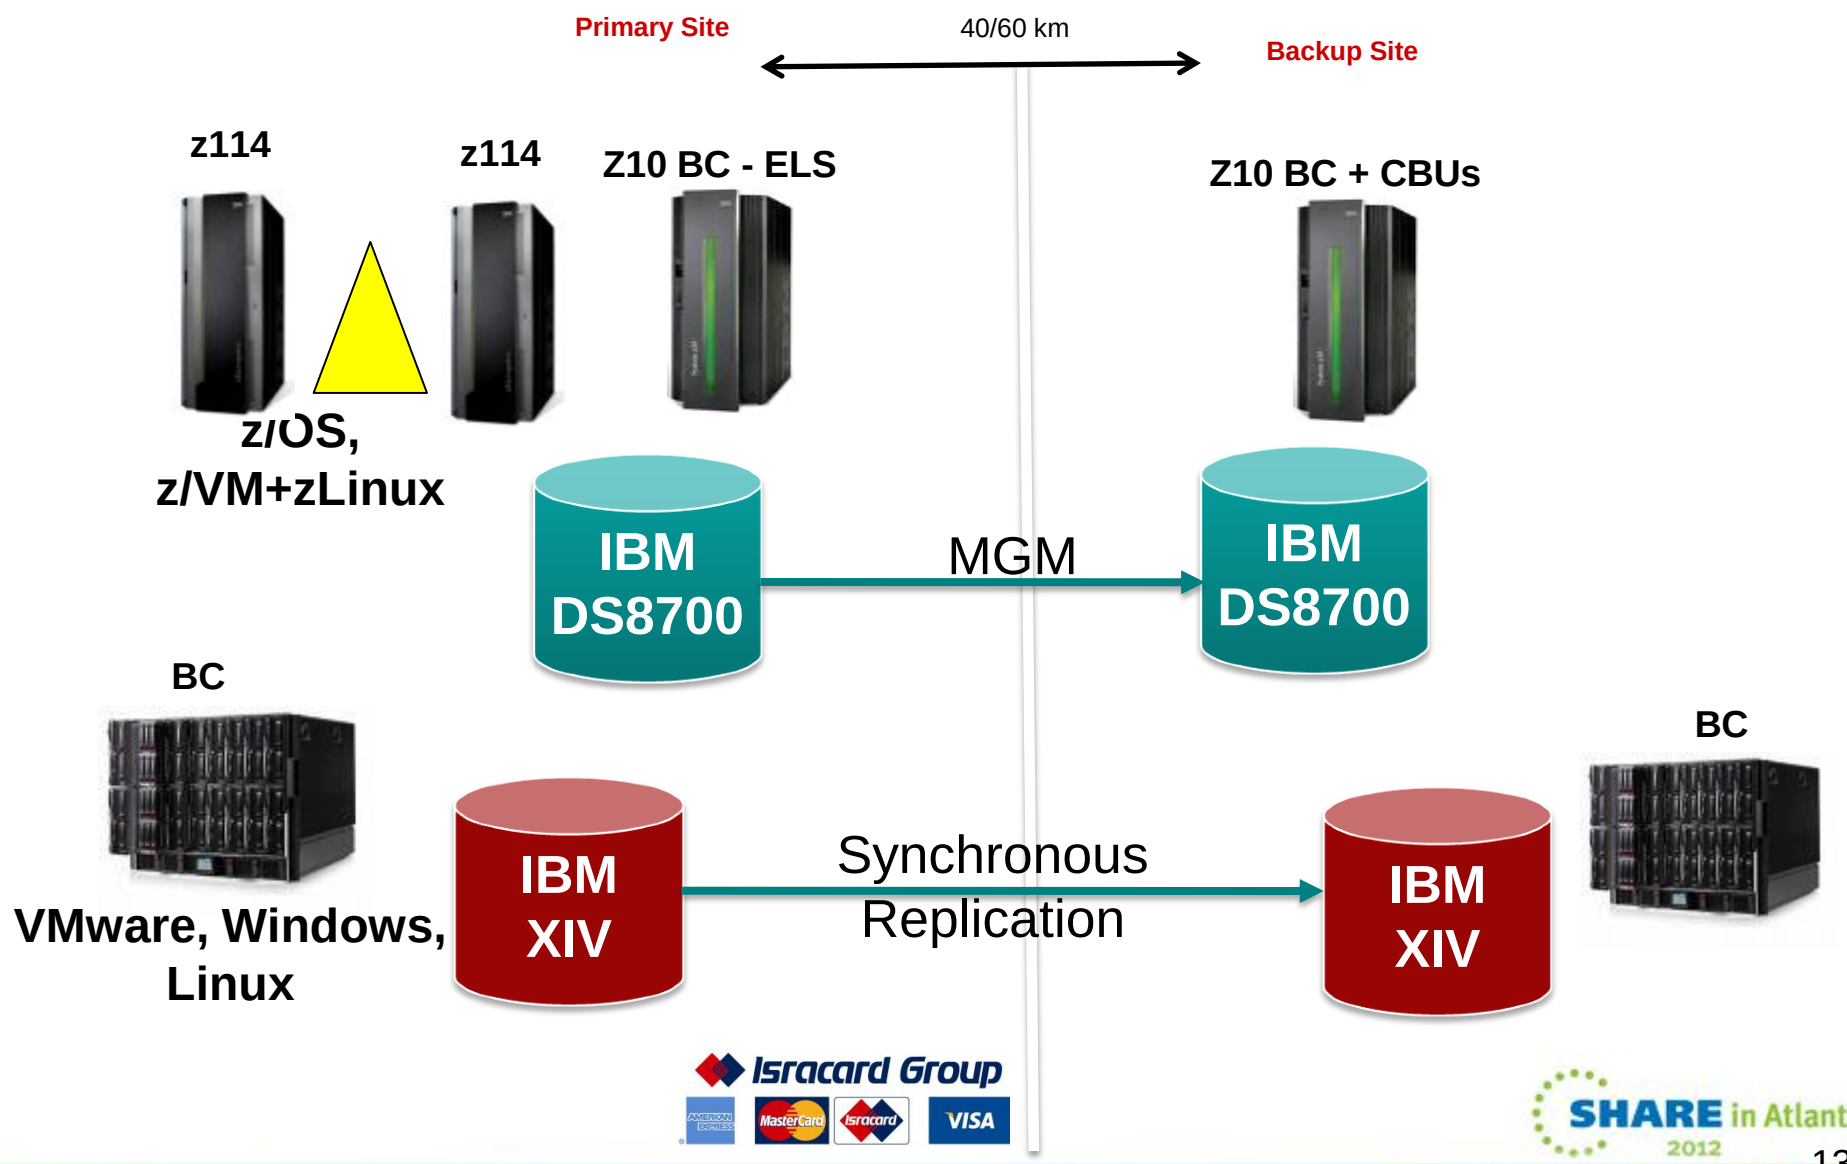

Agenda

- Introduction
- Some History
- **Current Status**
- Challenges and Solutions
- Observations
- Questions

Isracard zVM/zlinux Infrastructure

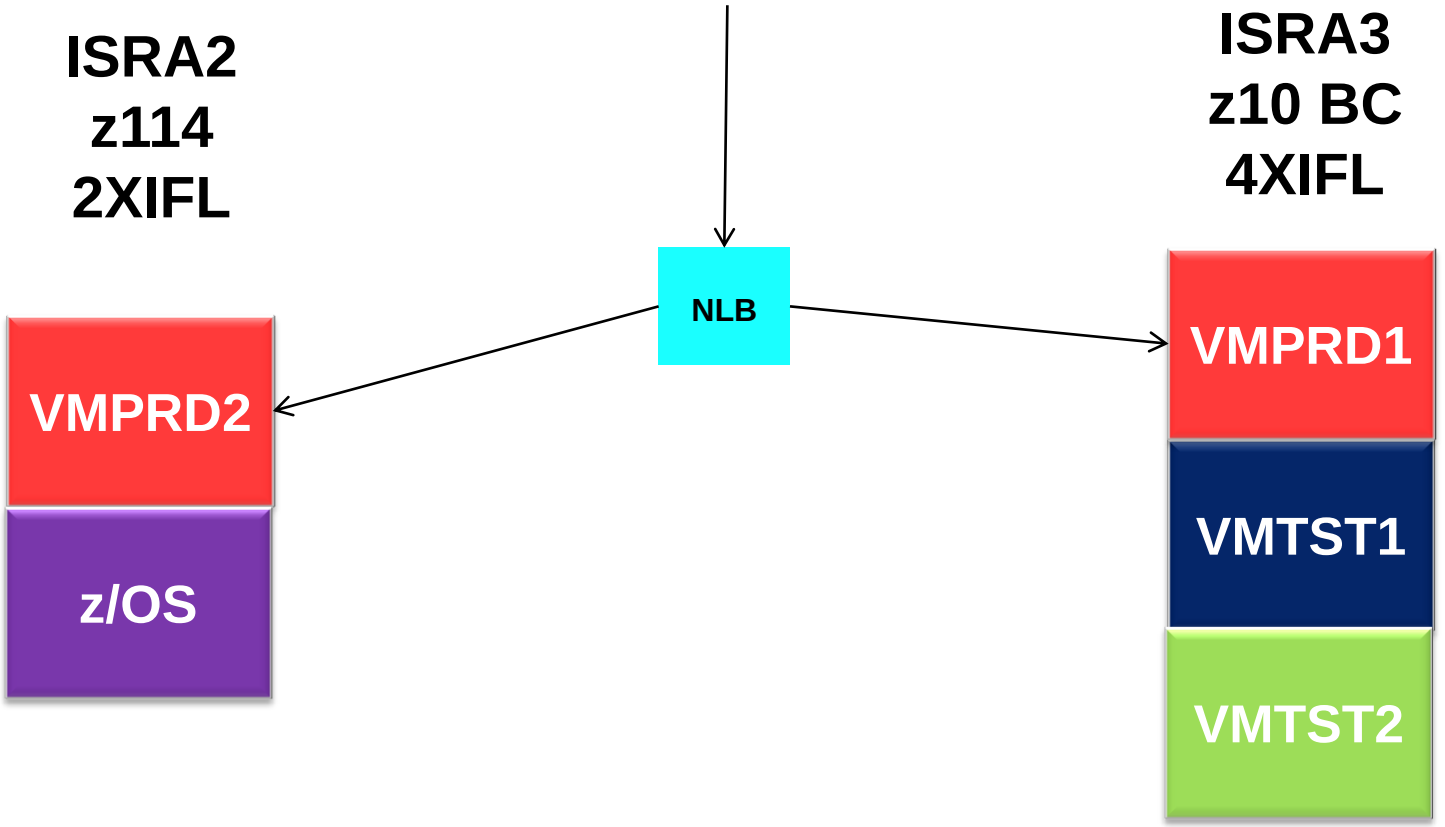

Current Status(1/3)

- 2 IFLS on a z114, 4 IFLs on a z10 BC
- Storage: DS8700/ECKD for binaries, XIV/FCP for data
- zVM 5.4 (6.2 planned for the summer)
  - 2 production LPARs
  - 2 test LPARs
  - 2 QA LPARs
  - 2 system LPARs
- RHEL 5.6 (6 planned for the end of 2012)
  - 110 linux images
  - 40 production

Current Status(2/3)

- Software
  - WAS
  - Oracle
  - WMB
  - iLog
- Applications
  - Internet site
  - Check Authorization
  - ESB
  - Rule Engine
  - FileNet

Current Status(3/3)

- Tools and utilities
  - BMC/Control-M(scheduling)
  - Symantec/Netbackup(backups)
  - CSL-Wave (managment and provisioning)
  - Omegamon for zVM(system performance)
  - CA Wily Introscope(application performance)
  - Tivoli System Automation(clustering)

Enterprise Linux Sever (ELS(

- z Series with IFLs only - specially priced
- The z10 BC and the z114 have only 10 engines
- 2CPs + 2 ICFs + 1 ZIIP + 2 IFLs = 7 CPUS
- What about growth and CBU/CoD?
- **Good:** No ELS at DR. We use CBUs on existing z10
- **Bad:** no Hipersockets(is this really bad? More later…)

CSL - WAVE

A provisioning tool

- Clone new images
- Allocate resources (disks, network interfaces, memory)

A management tool

- Activate/ Deactivate images
- Access (even if no network)
- Reports
- Automation

Basic Health checking

- CPU utilization
- Disk space running out

CSL - WAVE (1/4)

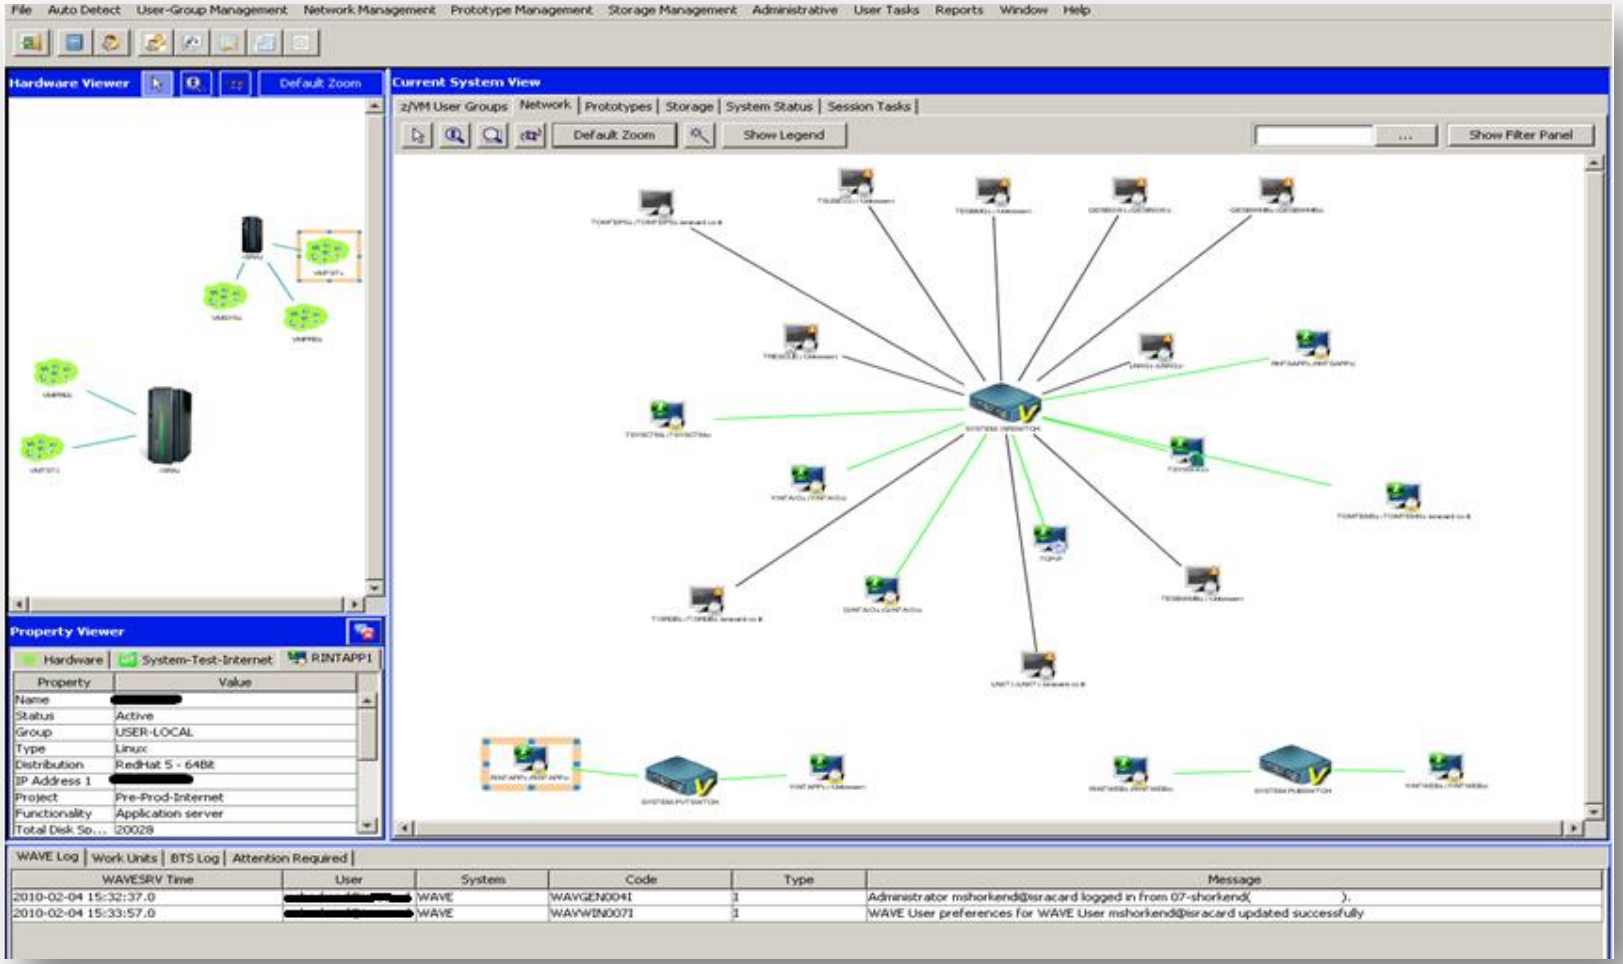
File Auto Detect User-Group Management Network Management Prototype Management Storage Management Administrative User Tasks Reports Window Help
Hardware Viewer
Default Zoom
Current System View
z/VM User Groups Network Prototypes Storage System Status Session Tasks
Default Zoom
Show Legend
Show Filter Panel
Property Viewer
Hardware System-Test-Internet R3NTAPP1
Property Value
Name
Status Active
Group USER-LOCAL
Type Linux
Distribution RedHat 5 - 64Bit
IP Address 1
Project Pre-Prod-Internet
Functionality Application server
Total Disk Sp... 20028
WAVE Log Work Units BTS Log Attention Required
WAVESRV Time User System Code Type Message
2010-02-04 15:32:37.0 WAVE WAVGEN0041 1 Administrator mshorkend@isracard logged in from 07-shorkend( ).
2010-02-04 15:33:57.0 WAVE WAVWIN0071 1 WAVE User preferences for WAVE User mshorkend@isracard updated successfully

CSL - WAVE (2/4)

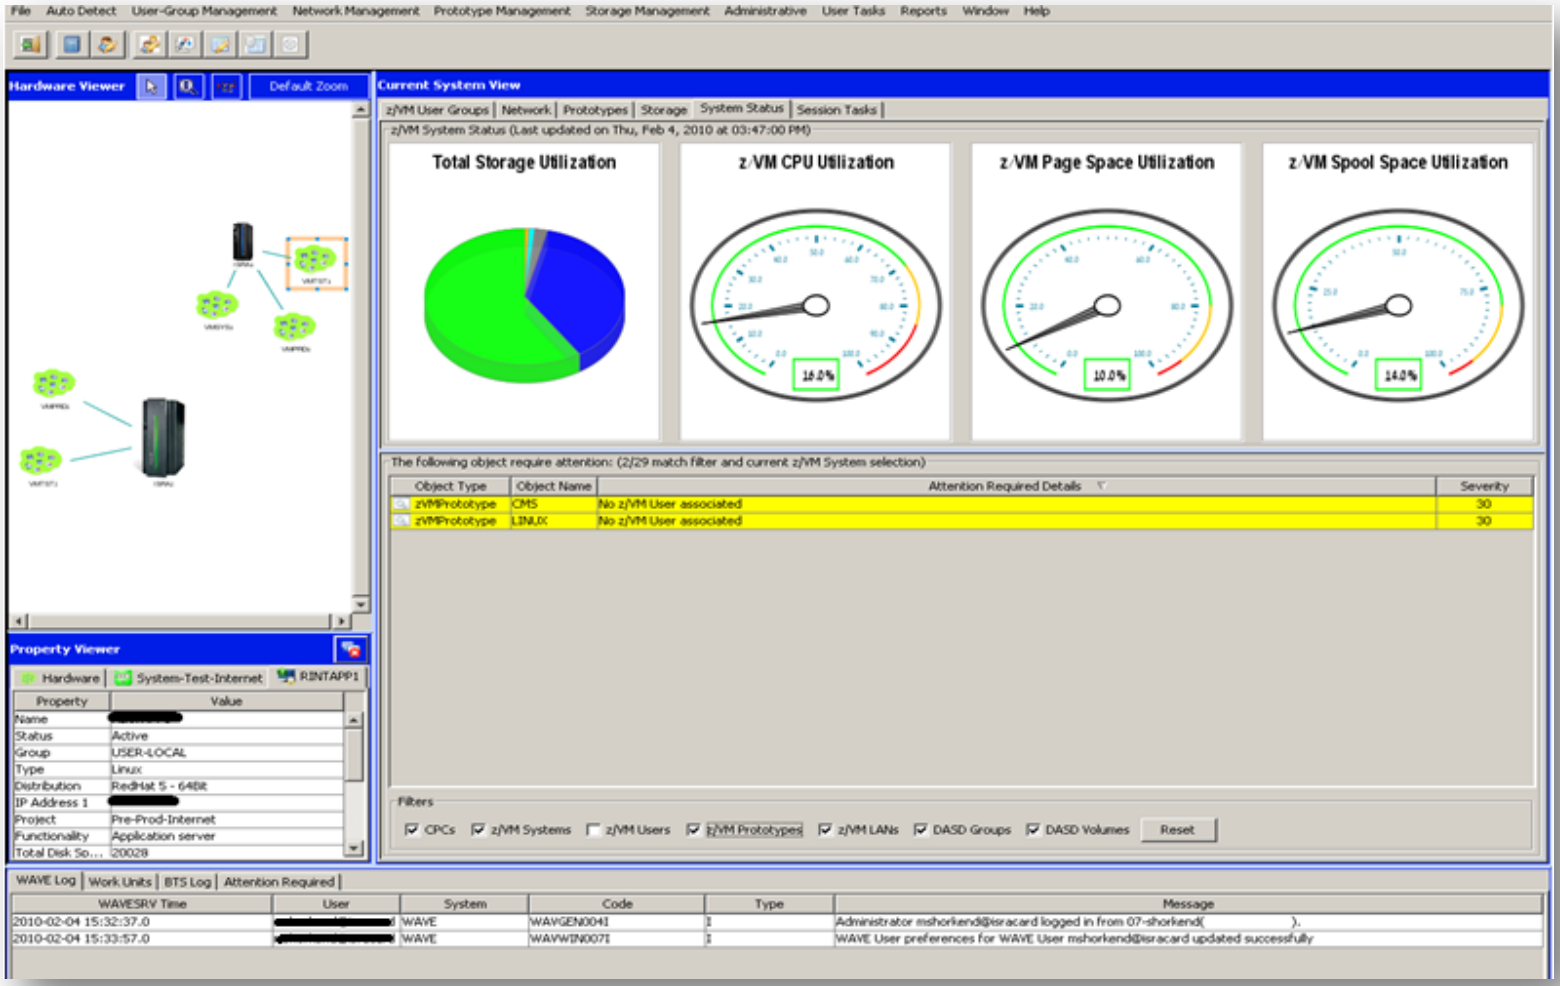
File Auto Detect User-Group Management Network Management Prototype Management Storage Management Administrative User Tasks Reports Window Help
Hardware Viewer
Default Zoom
Current System View
z/VM User Groups | Network | Prototypes | Storage | System Status | Session Tasks
z/VM System Status (Last updated on Thu, Feb 4, 2010 at 03:47:00 PM)
Total Storage Utilization
z/VM CPU Utilization
z/VM Page Space Utilization
z/VM Spool Space Utilization
16.0%
10.0%
14.0%
Property Viewer
Hardware | System-Test-Internet | RINTAPP1
Property Value
Name
Status Active
Group USER-LOCAL
Type Linux
Distribution RedHat 5 - 64Bit
IP Address 1
Project Pre-Prod-Internet
Functionality Application server
Total Disk Sp... 20028
The following object require attention: (2/29 match filter and current z/VM System selection)
Object Type Object Name Attention Required Details Severity
zVMPrototype CMS No z/VM User associated 30
zVMPrototype LINUX No z/VM User associated 30
Filters
CPCs z/VM Systems z/VM Users z/VM Prototypes z/VM LANs DASD Groups DASD Volumes Reset
WAVE Log | Work Units | BTS Log | Attention Required
WAVESRV Time User System Code Type Message
2010-02-04 15:32:37.0 WAVE WAVGEN0041 I Administrator mshorkend@isracard logged in from 07-shorkend( ).
2010-02-04 15:33:57.0 WAVE WAVWIN0071 I WAVE User preferences for WAVE User mshorkend@isracard updated successfully

CSL - WAVE (3/4)

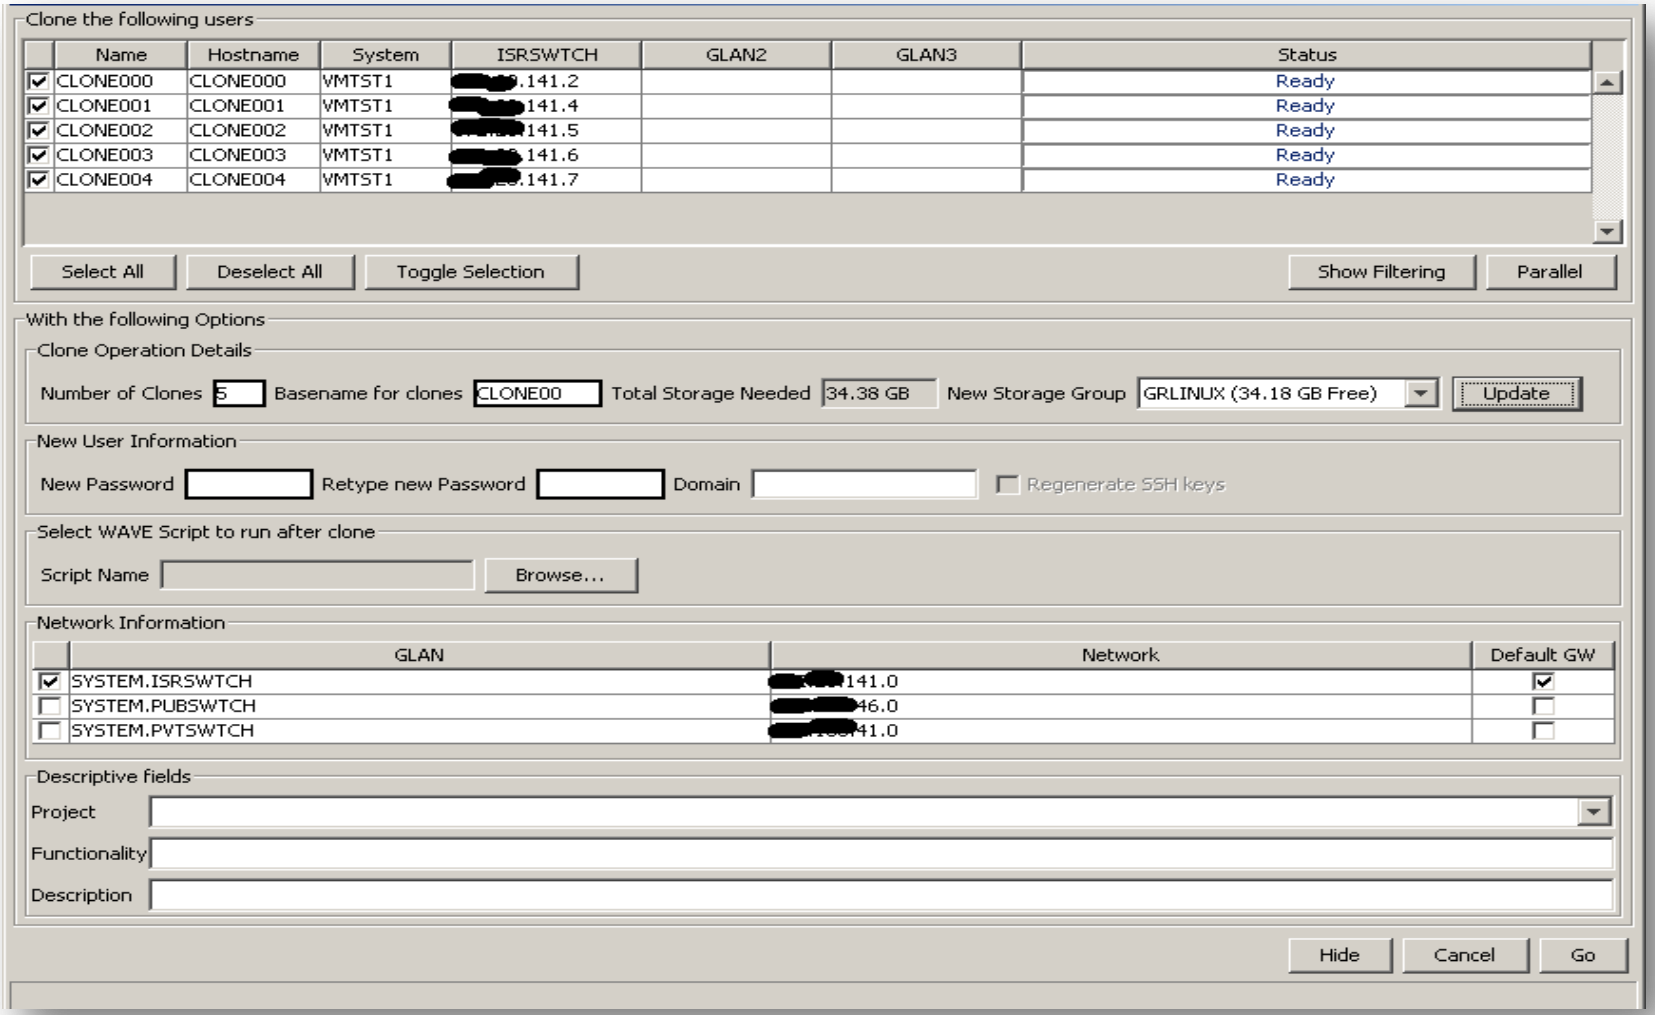
Clone the following users
Name
Hostname
System
ISRSWTCH
GLAN2
GLAN3
Status
CLONE000 CLONE000 VMTST1 .141.2 Ready
CLONE001 CLONE001 VMTST1 141.4 Ready
CLONE002 CLONE002 VMTST1 141.5 Ready
CLONE003 CLONE003 VMTST1 141.6 Ready
CLONE004 CLONE004 VMTST1 141.7 Ready
Select All
Deselect All
Toggle Selection
Show Filtering
Parallel
With the following Options
Clone Operation Details
Number of Clones 5
Basename for clones CLONE00
Total Storage Needed 34.38 GB
New Storage Group GRLINUX (34.18 GB Free)
Update
New User Information
New Password
Retype new Password
Domain
Regenerate SSH keys
Select WAVE Script to run after clone
Script Name
Browse...
Network Information
GLAN
Network
Default GW
SYSTEM.ISRSWTCH 141.0
SYSTEM.PUBSWTCH 46.0
SYSTEM.PVTSWTCH 41.0
Descriptive fields
Project
Functionality
Description
Hide
Cancel
Go

CSL - WAVE (4/4)

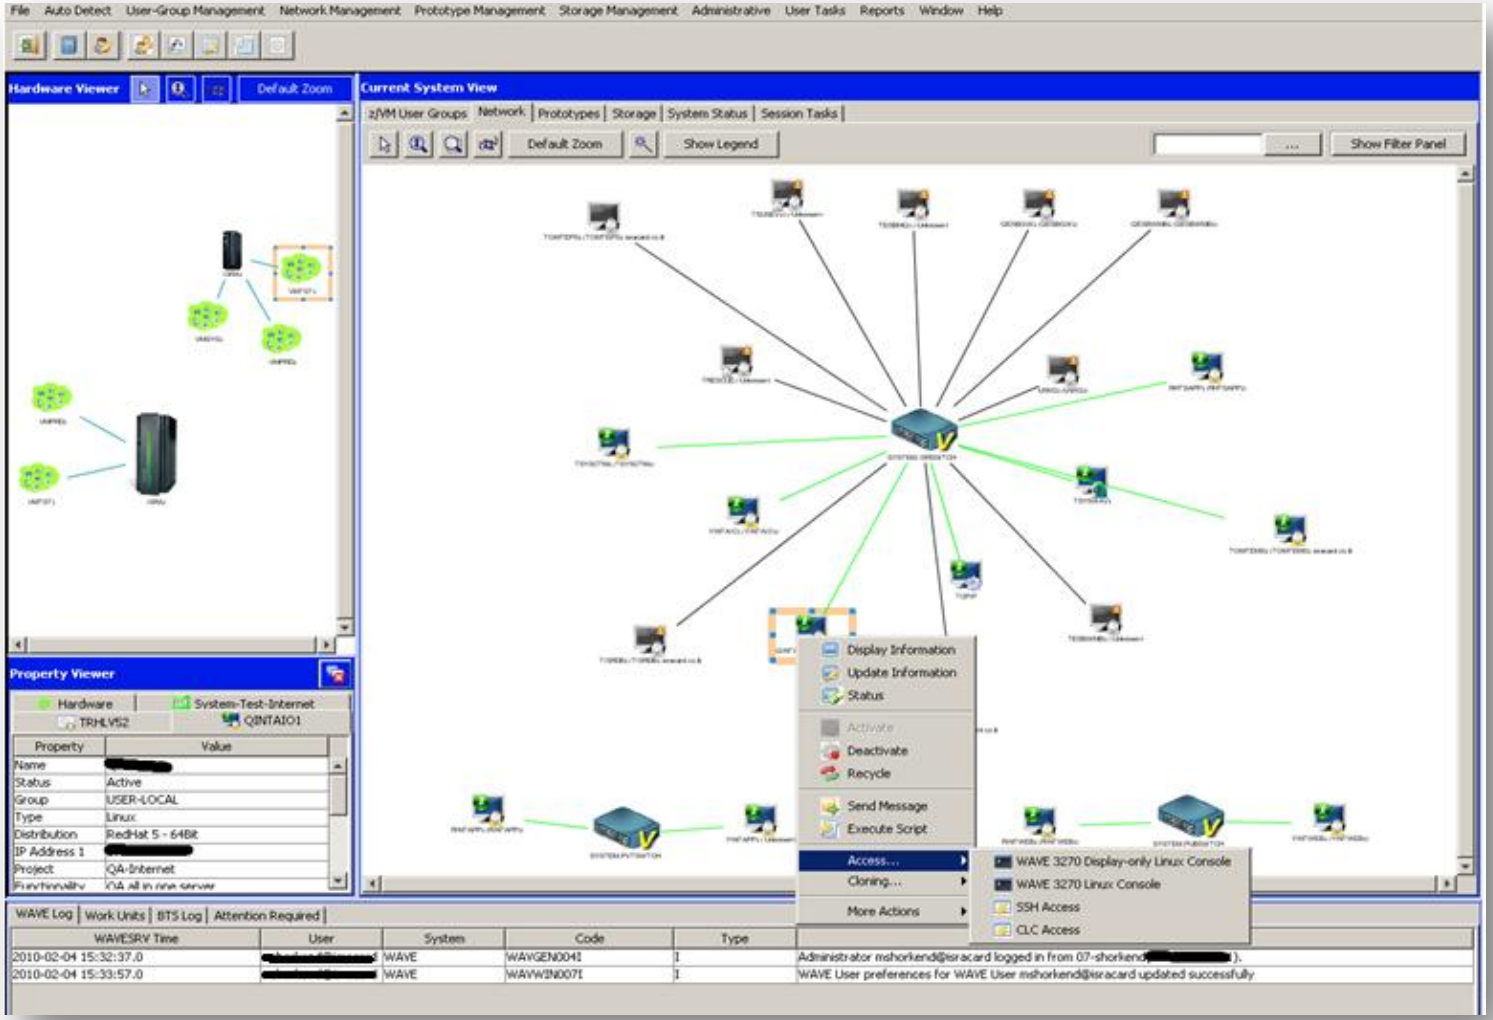
File Auto Detect User-Group Management Network Management Prototype Management Storage Management Administrative User Tasks Reports Window Help
Hardware Viewer
Default Zoom
Current System View
z/VM User Groups Network Prototypes Storage System Status Session Tasks
Default Zoom
Show Legend
Show Filter Panel
Property Viewer
Hardware
System-Test-Internet
TRHLV52
QINTAIO1
Property
Value
Name
Status
Active
Group
USER-LOCAL
Type
Linux
Distribution
RedHat 5 - 64Bit
IP Address 1
Project
QA-Internet
Display Information
Update Information
Status
Activate
Deactivate
Recycle
Send Message
Execute Script
Access...
Cloning...
More Actions
WAVE 3270 Display-only Linux Console
WAVE 3270 Linux Console
SSH Access
CLC Access
WAVE Log
Work Units
BTS Log
Attention Required
WAVESRV Time
User
System
Code
Type
2010-02-04 15:32:37.0
WAVE
WAVGEN0041
I
Administrator mshorkend@isracard logged in from 07-shorkend
2010-02-04 15:33:57.0
WAVE
WAVWIN0071
I
WAVE User preferences for WAVE User mshorkend@isracard updated successfully

TEP and Omegamon/VM (1/2)

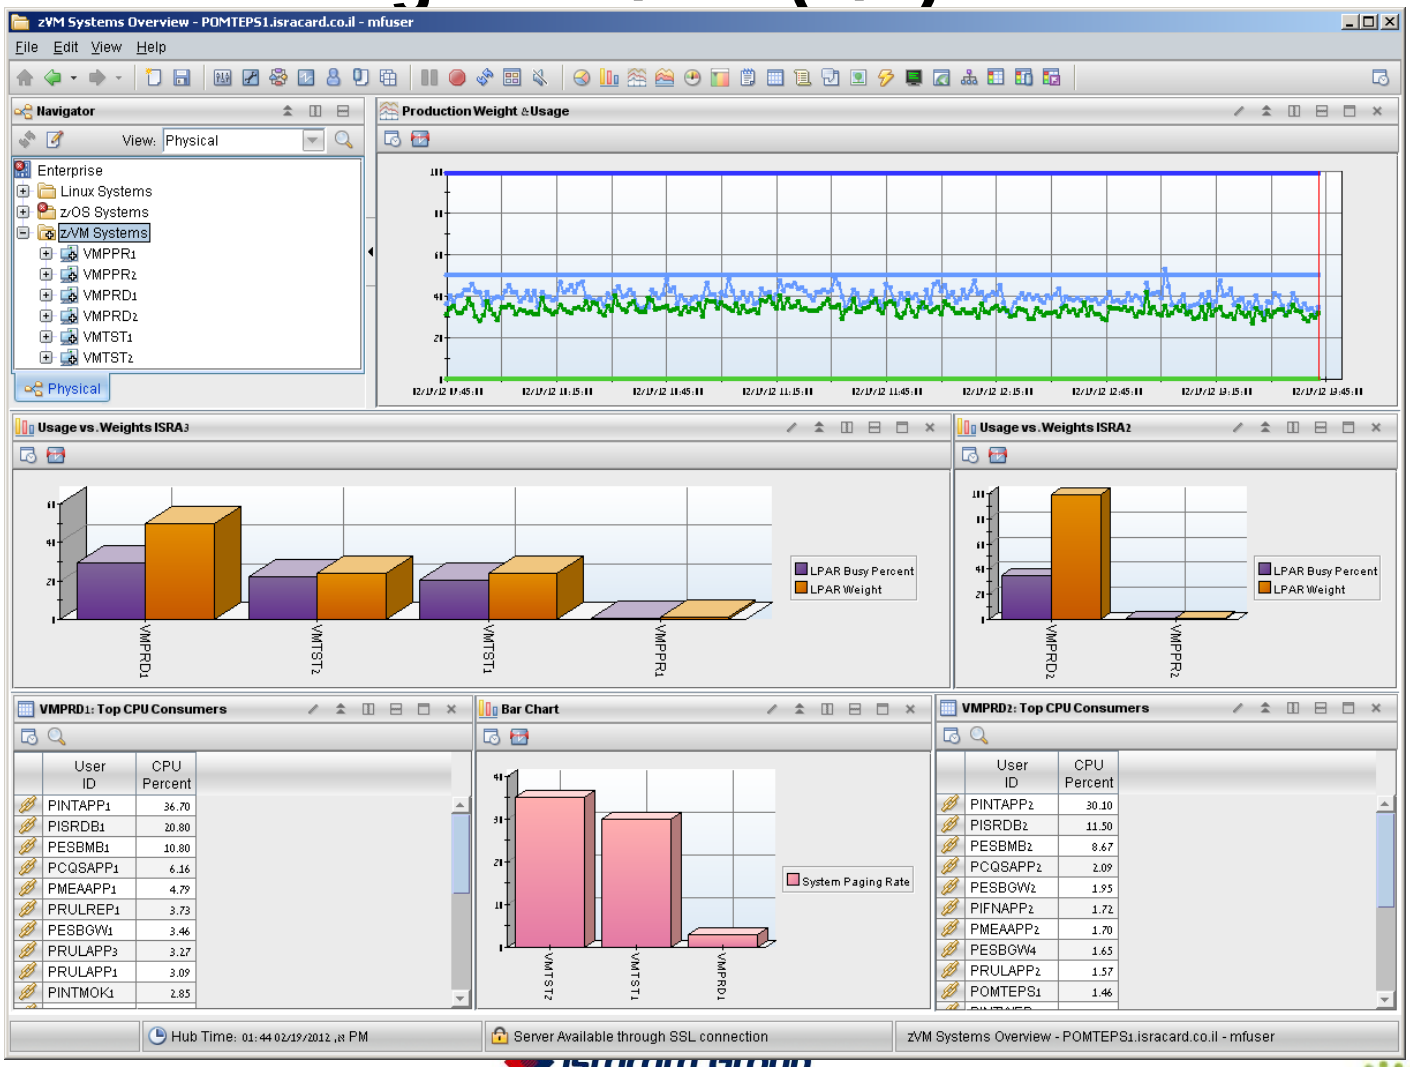

TEP and Omegamon/VM (2/2)

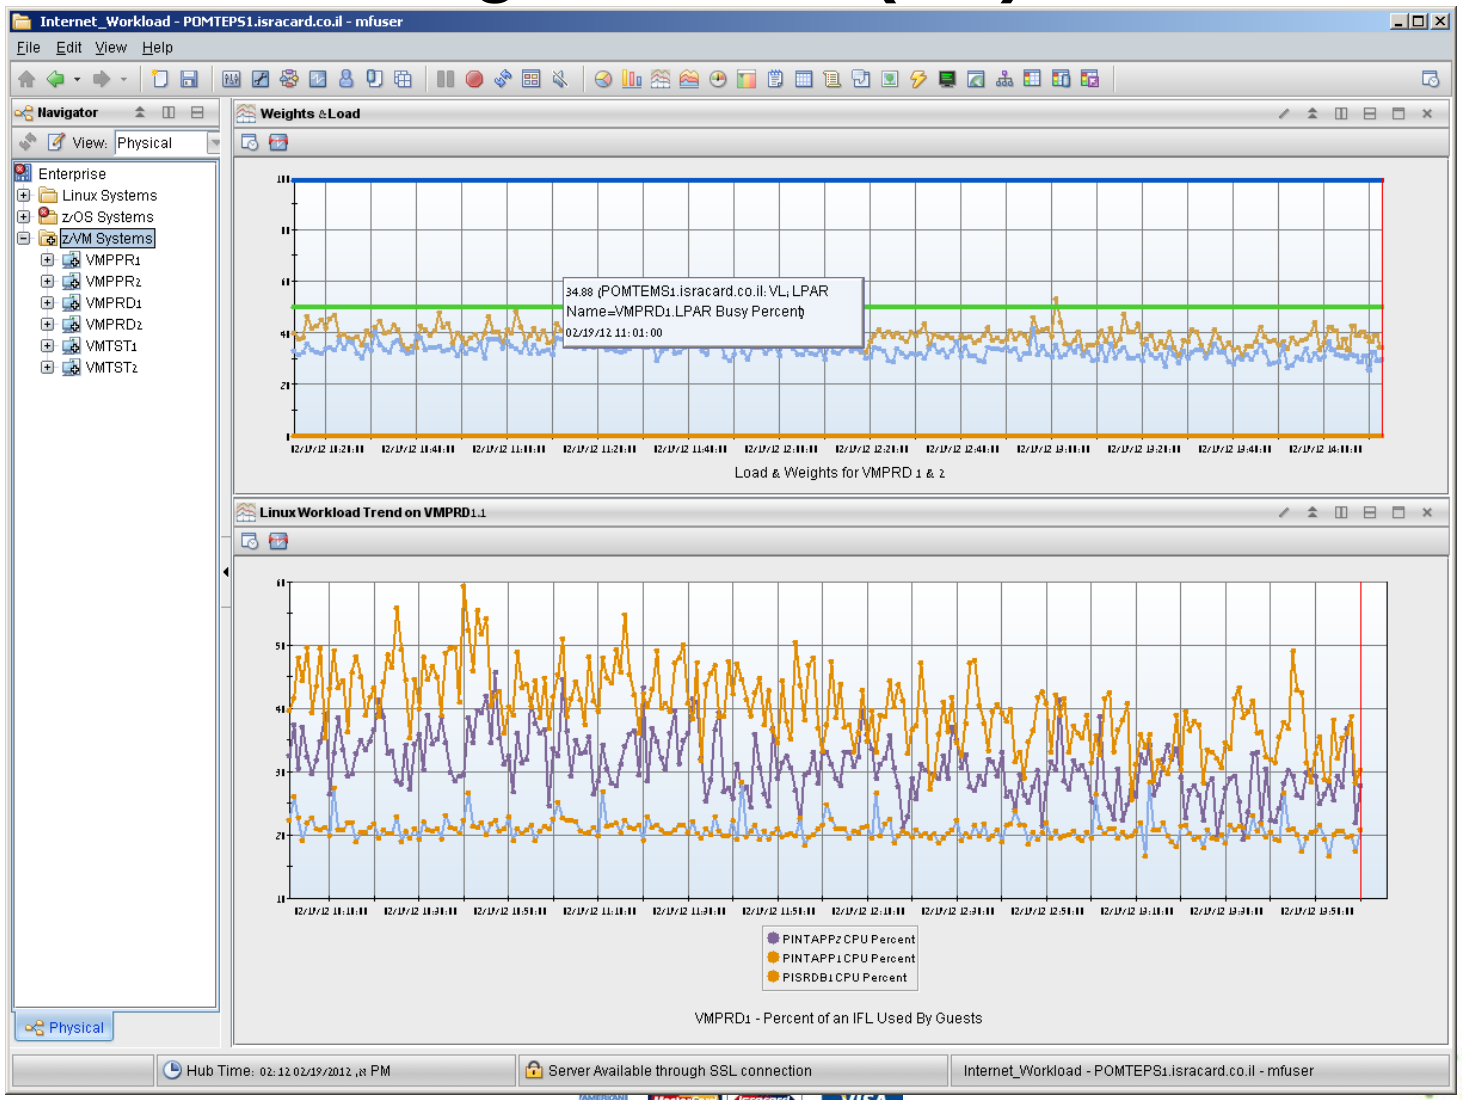

Agenda

- Introduction
- Some History
- Current Status
- **Challenges and Solutions**
- Observations
- Questions

The Challenges and some answers(1/2)

- **The newer Intel cores are impressive**
  - *These cores present a very viable alternative to z10 BC*
  - ✓ *Our early experience with z114 is very promising*
- **zVM overhead is excessive**
  - *The workload per IFL is below expectations*
  - ✓ *We are working closely with IBM to resolve this issue*
  - ✓ *Migration to zVM 6.2 should help*
- **Some applications work better on Intel**
  - *Our previous strategy was to put everything on z.*
  - ✓ *Adjust policy to 'fit for purpose'*

The Challenges and some answers(2/2)

- **Some software is not certified for RHEL/z(or gets certified later)**
  - *Specifically: Clustering software, antivirus, Oracle, DBA Monitoring tools, Asset Management,C² software*
  - ✓ *We are working closely with IBM and RHEL to alleviate this*
- **Performance perceptions**
  - *Users (end users, sysadmins, DBAs) do not like to share*
  - ✓ *Tuning. And we try to listen to our users.*

Agenda

- Introduction
- Some History
- Current Status
- Challenges and Solutions
- **Some Observations**
- Questions

Some general observations (1/2)

- **Different versions of RH for different software**
  - *would you keep z/OS 1.9 for DB2 8 and z/OS 1.11 for CICS/TS 4.1?*
- **Bleeding edge at times**
  - *Certification - not always there*
    - *Oracle 11g certification came in too late*
  - *Sometime we had to wait for software to be written*
  - *Not all software is supported on z*
- **Hipersockets – we have not found a justification for it (yet)**

Some general observations (2/2)

- **Managerial issues**
  - *Is it Mainframe or Distributed? -* ***Try to avoid turf wars!***
    - *We decided to manage it in one place*
  - *You need a full time z/VM expert*
  - *DBAs and Sysadmins do not like virtual platforms -* ***Educate, Educate, Educate***
- **Business Class Issues**
  - *Processor power -* ***Most TCO studies were performed for EC***
  - *Total number of CPUs = 10 -* ***Forced us to go to ELS (Enterprise Linux Server)***

They multiply

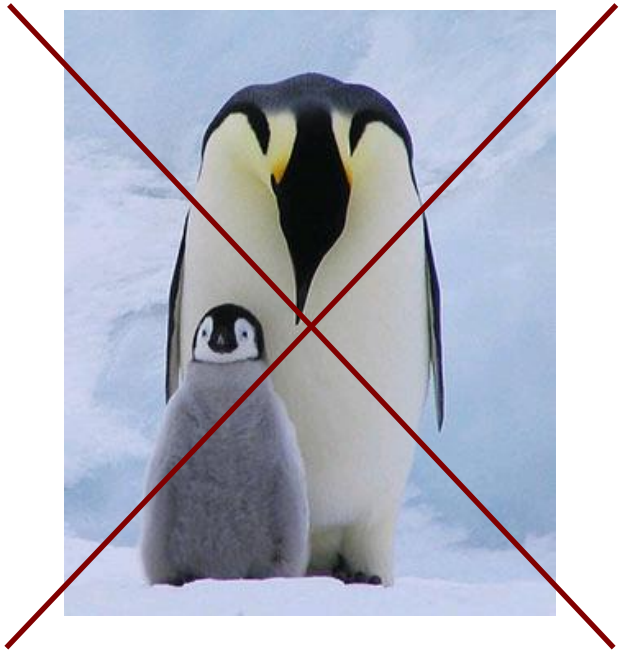

Agenda

- Introduction
- Why (z)Linux?
- Chronological road to production
- Some tools
- Observations
- **Questions**

Questions ?