

z/VM and Linux Disaster Recovery – A Customer Experience

Lee Stewart
Sirius Computer Solutions (DSP)

Date Thursday, August 14th, 2008
Session 9210

The Business Partner

- **Sirius Computer Solutions**
 - No, not the satellite radio people




- **IBM Reseller**
 - Not a DR vendor



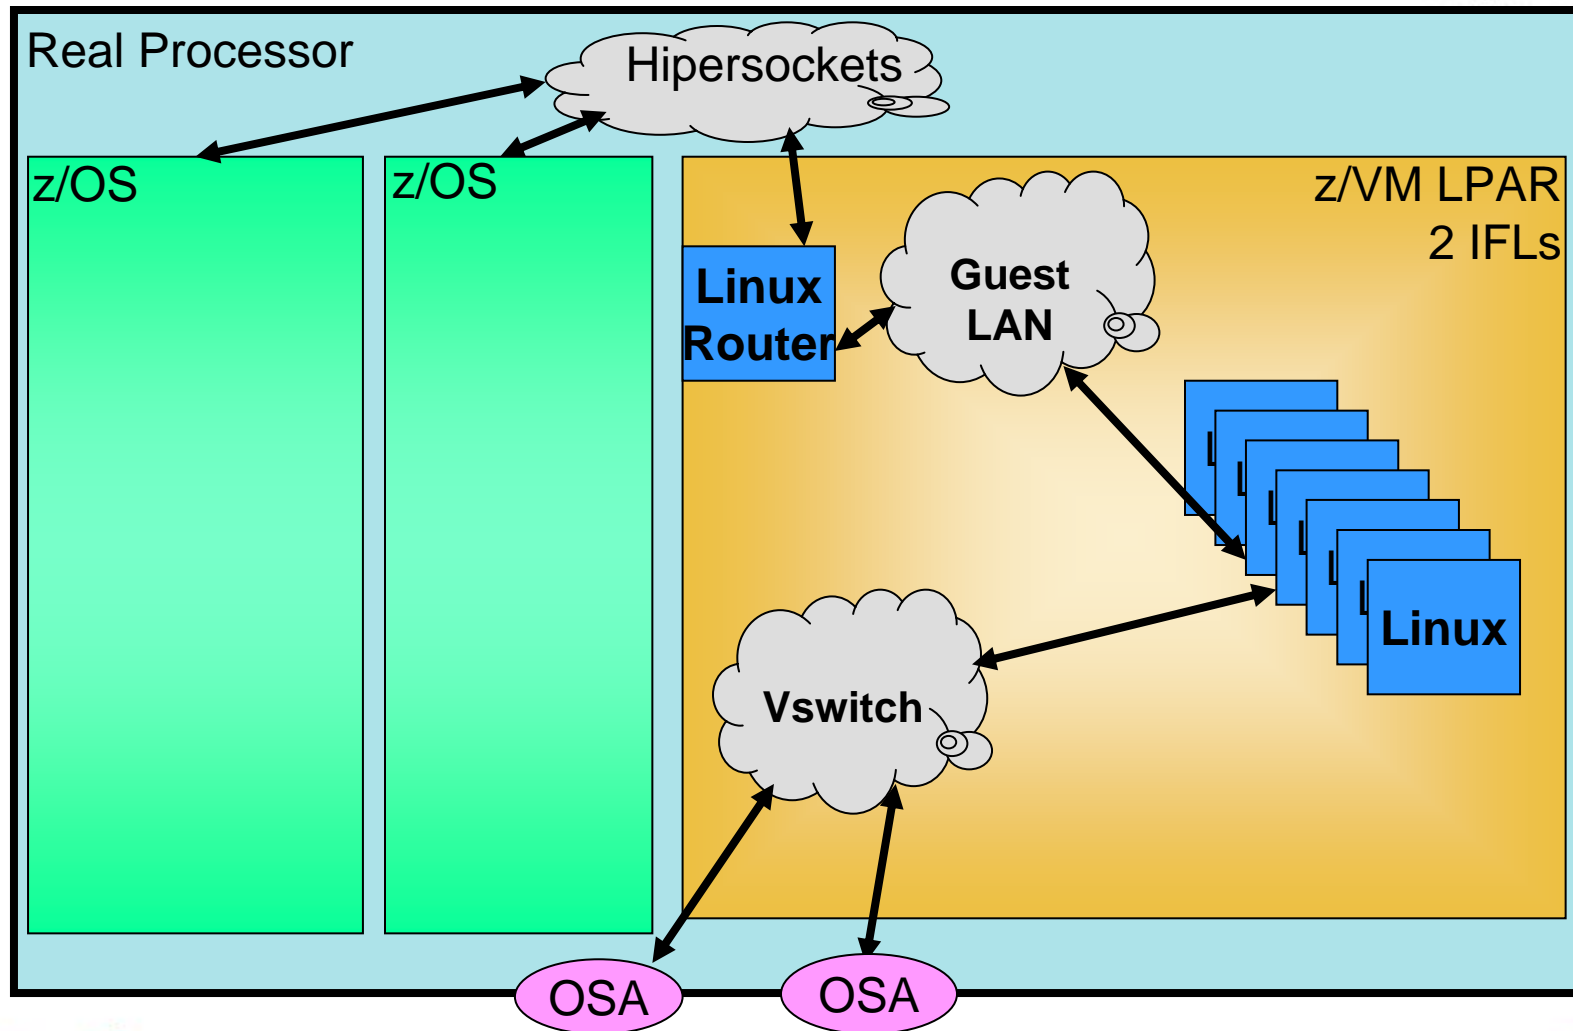
- **Most hardware sales are bundled with software services**
- **My role**

The Customer

- **“An Insurance Company”**
- **z9 BC with 2 IFLs**
- **2 Production z/OS LPARs**
- **1 VM & Linux LPAR**
- **z/VM 5.2; SLES9 SP3; 20+ Linux Images**
 - 8-10 “Production”
 - Migrating to SLES10 SP2
- **WebSphere; MQ; DB2 Connect**
- **Fair Isaac – Blaze Advisor**



The Customer Setup



Two Part Backups

- **Full pack backup from z/OS weekly**
 - Linux machines shut down briefly on Sunday evenings, FLASHCOPYed, then backed up to tape.
 - z/VM backed up periodically
- **TSM backup for incremental changes and file level restores**
 - Server on a Windows box

The DR Vendor


- **“A Large DR Vendor”**
- **Large z9**
 - **Plenty of Memory**
 - **Lots of Processors**
 - **All CPs, no IFLs**
 - **All “full speed”**
- **z/OS DR “always done under VM in the past”**

The Previous DR Test

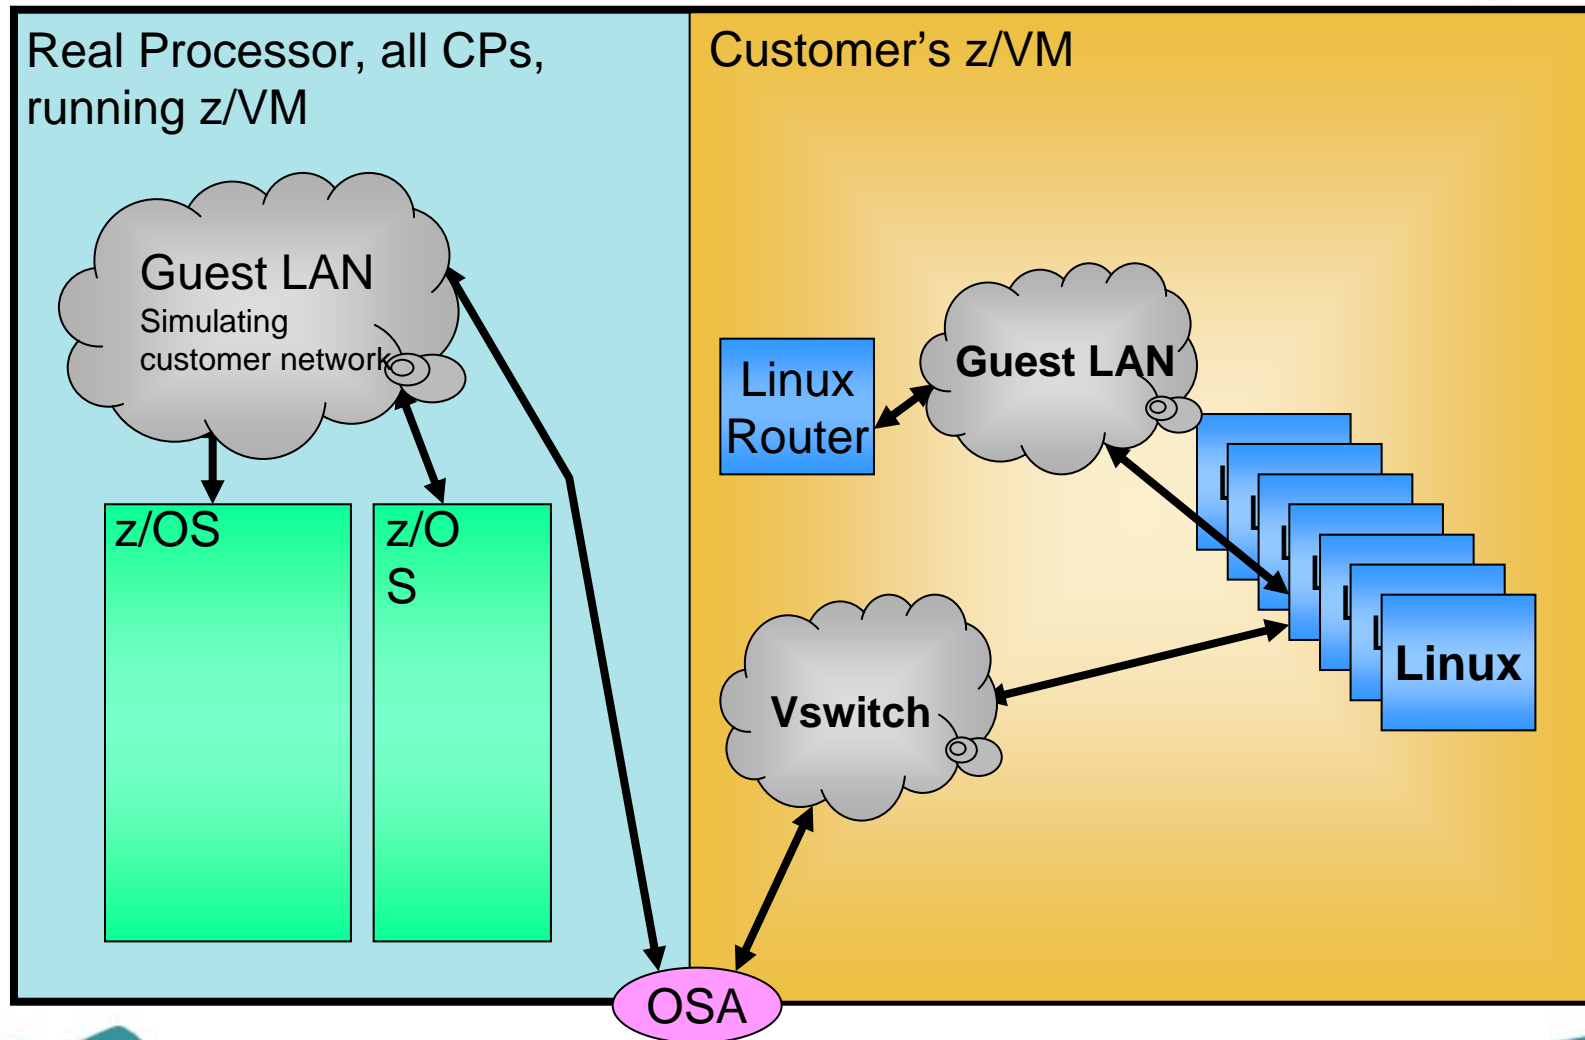
- **A year earlier, before z/VM and Linux**
- **Started on a Monday morning**
- **z/OS LPARs Restored, Up and Running by Monday afternoon**
- **“Distributed Systems” “Mostly up” by the end of Wednesday**
 - **In fairness, a large hodgepodge of systems**
 - **Some never made it up in full 4 day test**
- **Successful test, but “next time needs to be better”**

Last Year's Test

- **Fall '07**
 - Monday to Thursday
 - Customer personnel NOT onsite at DR vendor site
- **2 z/OS LPARs**
- **1 z/VM LPAR**
- **20+ Linux Guests**
 - 8-10 Production
 - “Why not bring them all up”
- **Still many “Distributed Systems”**
 - And many moved to Linux on Z
 - Fewer unique environments



Last Year's DR Setup



Last Year's Test

- **Monday morning:**
 - Start restoring tapes
- **Monday early afternoon:**
 - 2 z/OS LPARs – Up and running
 - 1 z/VM LPAR – Up and running
 - 20+ Linux Guests – All Starting
 - Distributed Systems – “Working on it”

First Problem




- **Monday late afternoon**
- **Performance!**
- **Customer's VM image only given less than ½ the requested memory, and no Xstor**
- **VM Paging at thousands of pages per second**
- **First fix attempt – force off “non-production” Linux guests**
- **Still slow, still heavy paging**
- **Shutdown VM, increase VM's memory, restart**


Second Problem



- **Tuesday mid-morning**
- **Performance!**
- **Memory is now correct, even Xstor**
- **VM can now see 2 real processors but begin to suspect they aren't dedicated to our "LPAR"**
- **Discover VM is NOT in it's own LPAR, but under the DR vendor VM!**
- **Can't shutdown z/OS systems to reconfigure LPARs**



Last Year's Real DR Setup



Second Problem – Why?

- **Customer VM & CMS runs great, Linux runs terrible**
- **Why? Hardware vs. Software**
- **System Z hardware only supports 2 levels of virtualization – 2 levels of SIE**
- **LPAR support uses the 1st level SIE to run the LPARs (DR Vendor's VM system)**
- **The DR vendor's VM system (1st level) uses the 2nd level SIE to run the 2nd level systems (the customer's z/OS & VM images)**
- **The customer's VM system (2nd level) cannot use SIE to manage the Linux guests (3rd level) – all privops and management tasks must be simulated by CP**

Second Problem – The Fix


- **Avoid that 2nd level of VM**
- **Two ways – depending on DR vendor and your configuration**
 - #1 - Put your VM in an LPAR on the DR vendor's machine**
 - Recreates more of your environment
 - Easiest if a large number of Linux guests
 - Change the OSA addresses for VM's TCPIP and Vswitches
 - Usually costs a little more
 - #2 - Run your Linux guests directly under the DR vendor's VM**
 - Your Vswitch and guest LANs have to be created by the DR vendor
 - You don't have "all your stuff" in case
 - Easiest if you are only going to run a couple Linux guests
 - Usually a little cheaper

Second Problem – Bypass

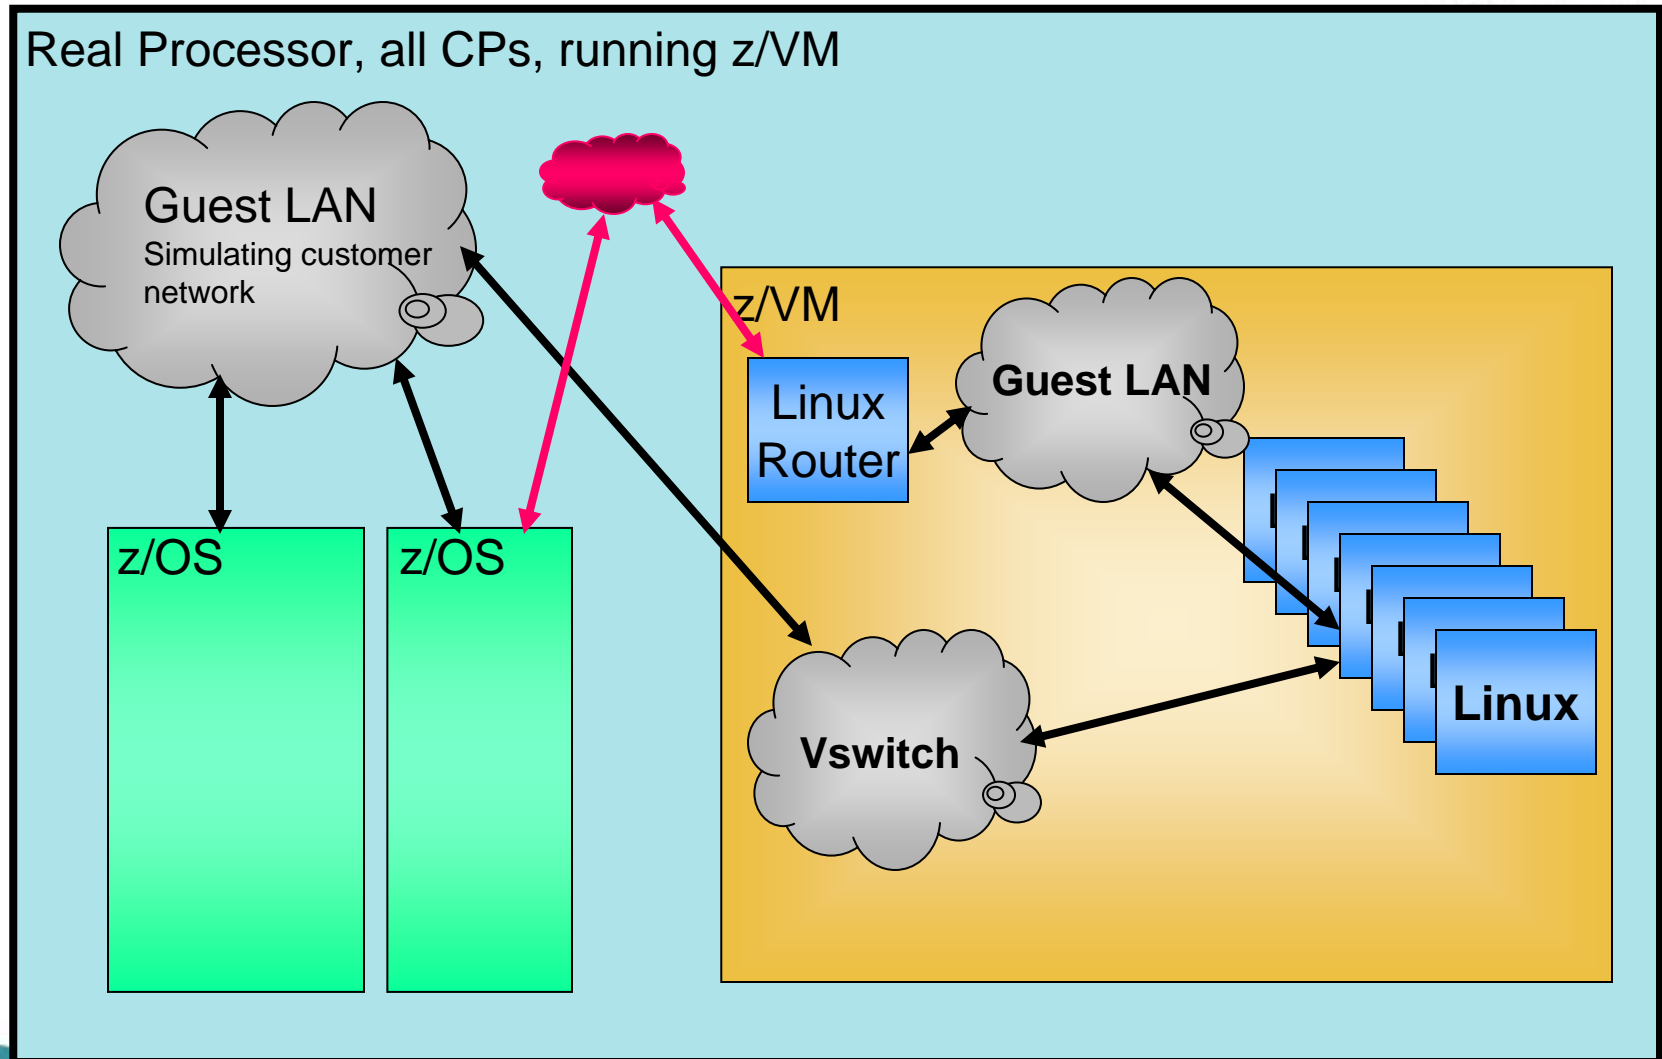
- **Customer decides bringing up a single production Linux will be a successful DR test – this time**
- **DR vendor dedicates 2 CPs to the VM userid**
- **Kill all but one Linux guest**
- **Slow but livable – for the test**

Third Problem

- **The Linux guest runs ok, but can't contact z/OS**
- **Oh yea... Start the Hipersocket router**
- **Slow, but it comes up**
- **Guest talks to router ok, but not to z/OS**
- **No DR definition for Hipersocket connection**




The Missing Hipersocket Router

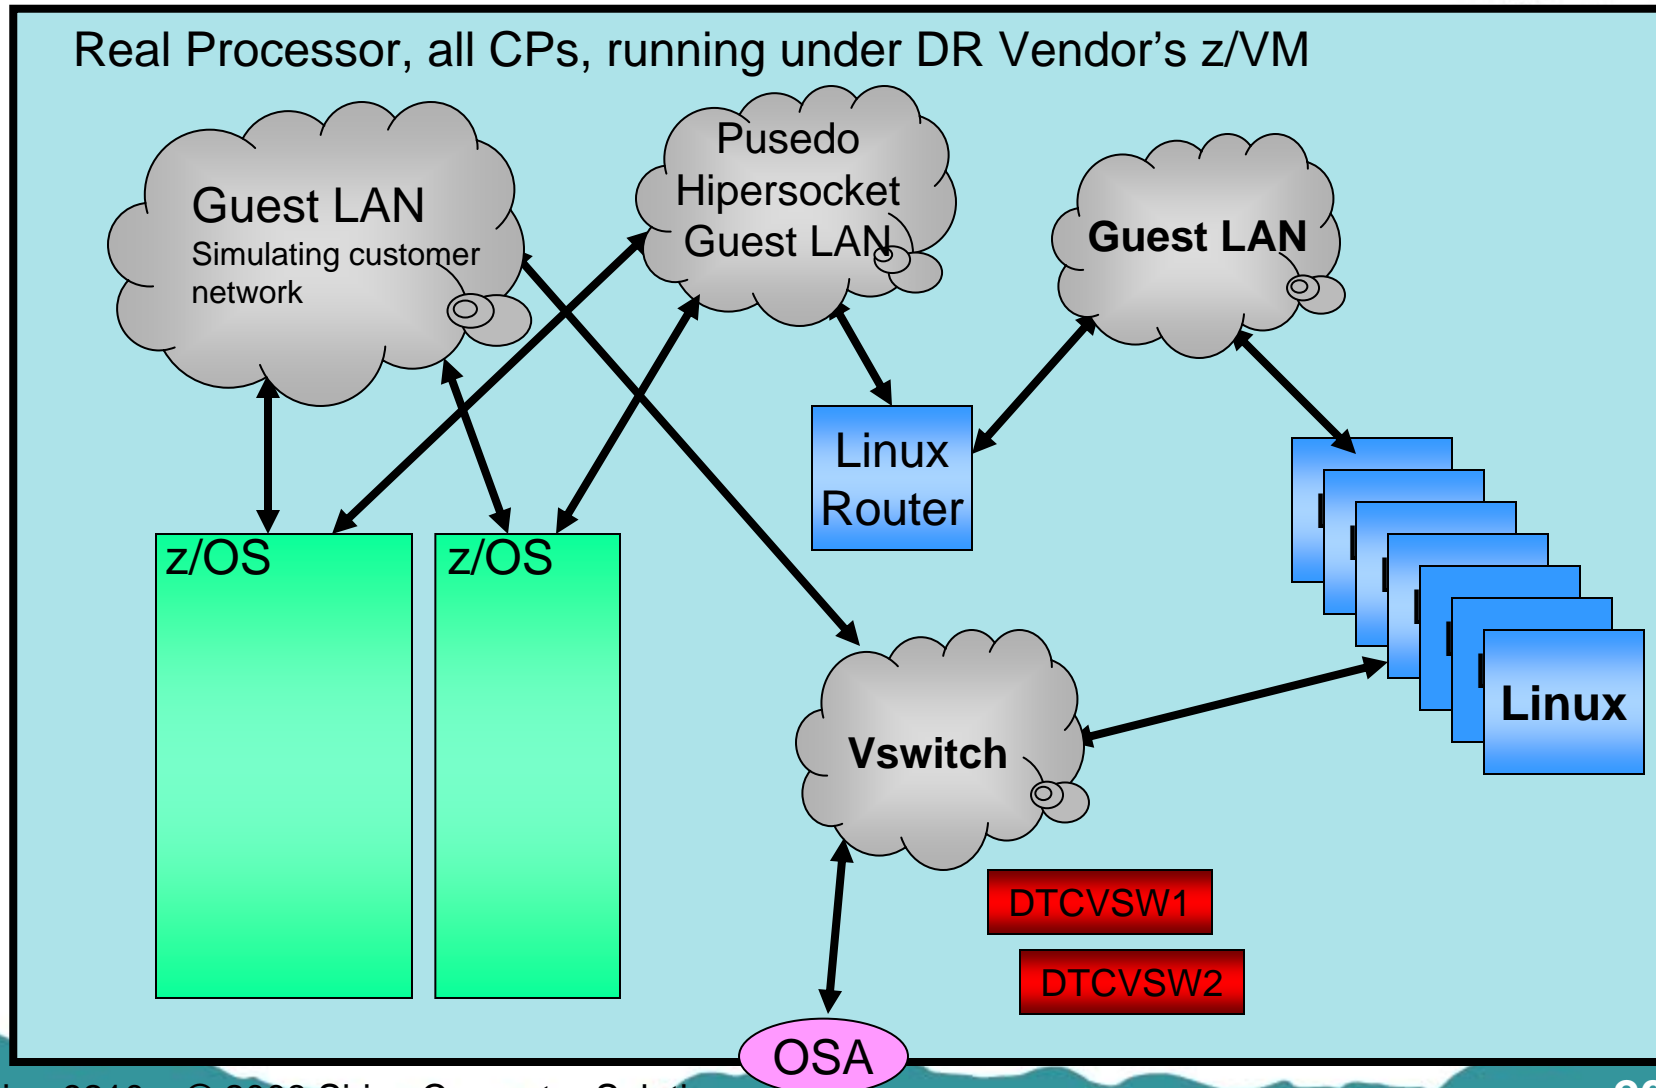


This Year's DR Test

- **Customer's Linux Under DR Vendor's VM**
 - Decided by DR coordinator – for cost reasons
- **Monday morning:**
 - Start restoring tapes
- **Monday early afternoon:**
 - 2 z/OS LPARs – Up and running
 - 3 Linux machines – Up and running (under the DR VM)
 - Only the production machines
 - Distributed Systems – “Working on it”
- **One “small” problem**
 - Missing Vswitch controller userids



This Year's Real DR Setup



Next Year's DR Test

- **Customer VM in a DR LPAR**
 - **Continue to migrate non-mainframe workload to Linux on z**
 - DR recovery speed is only one benefit
 - **Bring up all the Linux machines**
 - **DR their Development and Test environments as well**
 - **Expect all z/OS and z/VM and Linux on System z to be up and running by Monday afternoon.**

Lessons Learned

1. Communicate!

- What we thought was going to be the configuration, wasn't
- What we thought was being changed, wasn't
- Miscommunication all the way around – sysprog, DR coordinator, DR contract, DR sales, DR support

2. Be proactive

- Don't just tell your DR coordinator once and leave it all to them
- Don't skip out on DR planning meetings
- You know your systems much better than anyone else

Lessons Learned

3. Be SURE of your DR configuration

- For each LPAR
- Hardware – Processors, memory, OSAs
- Network – all connections, including Hipersockets
- Know what will be in LPARs and what will be under VM


4. Plan ahead to avoid 3rd level Linux

- No Linux under VM under VM
- Run your VM in a DR LPAR
- or -
- Run your Linux directly under the DR VM

Lessons Learned

- 5. Check your configuration when you get on the DR machine**
 - Hardware, memory, network
 - Don't assume it's set up right
 - Don't assume you have all the pieces you need


- 6. Don't forget Hipersockets**
 - And if running a Hipersocket router, don't forget that
 - Simulate via a Guest LAN or real OSA if necessary



S H A R E
Technology • Connections • Results

Summary





Thank You

Session 9210

Lee Stewart
lee.stewart@siriuscom.com