



## z/VM and Linux Disaster Recovery – A Customer Experience

Lee Stewart  
Sirius Computer Solutions (DSP)

Friday, August 28, 2009 - 9:30am  
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 – z/VM & Linux on System Z

© 2009 Sirius Computer Solutions 2

---

---

---

---

---

---

---

---

---

---

### What Not To Do



Our Disaster Recovery Plan Goes Something Like This...



© 2009 Sirius Computer Solutions 3

---

---

---

---

---

---


---

---

---

---

**The Story**



- **Two Companies**
  - An Insurance Company
  - A DR Vendor
- **Why no company names?**
- **Their DR test experiences over several years**
  - Before Linux on System Z
  - The first couple DR tests
    - What were the problems and how to avoid them
  - The latest DR tests
- **Told from the system programmer perspective**

© 2009 Sirius Computer Solutions 4

---

---

---

---

---

---


---

---

---

---

**Not Bashing the DR Vendor**



- Well respected DR vendor
- Some times Murphy wins
  - “What ever can go wrong, will go wrong.”
- Communication misunderstandings
  - Some inside the customer account
  - Some between the customer and DR vendor
- Story told from the system programmer perspective
  - Might appear differently if told by the DR vendor, or even the customer’s DR coordinator

© 2009 Sirius Computer Solutions 5

---

---

---

---

---

---


---

---

---

---

**The Customer**



- “An Insurance Company”
- z9 EC with 2 IFLs
  - Now a z10 EC with 3 IFLs
- 2 Production z/OS LPARs
- 1 VM & Linux LPAR
- z/VM 5.2; SLES9 SP3; 20+ Linux Images
  - 8-10 “Production” Linuxes
  - Now z/VM 5.4; SLES10 SP2
- WebSphere; MQ; DB2 Connect
- Fair Isaac – Blaze Advisor

© 2009 Sirius Computer Solutions 6

---

---

---

---

---

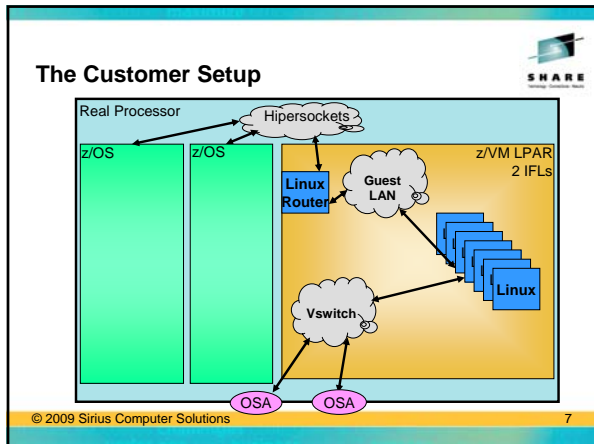
---

---

---

---

---




---

---

---

---

---

---

---

---

- ### Two Part Backups
- Full pack backup from z/OS weekly
    - All Linux apps require z/OS database access
    - z/OS and the databases are already down during a weekly backup window
    - Linux machines shut down briefly on Sunday evenings, FLASHCOPYed, then backed up to tape
    - Less than 10 minutes from when the Linuxes shutdown till they are all restarting
    - z/VM backed up periodically – doesn't change often
- © 2009 Sirius Computer Solutions 8

---

---

---

---

---

---

---

---

- ### Two Part Backups
- TSM backup for incremental changes and file level restores
    - Server on a Windows box
    - z/OS DR restore gives you a running Linux
      - TSM client already installed and configured
    - Minimal data to restore – in this client environment
  - Similar two part backup setup used at other clients with different backup software
- © 2009 Sirius Computer Solutions 9

---

---

---


---

---

---

---

---

**The DR Vendor** 

- “A Large DR Vendor”
- Large z9
  - Plenty of Memory
  - Lots of Processors
  - All CPs, no IFLs
  - All “full speed”
  - Now a z10
- z/OS DR “always done under VM in the past”

© 2009 Sirius Computer Solutions 10

---

---

---


---

---

---

---

---

**The Pre-Linux DR Test** 

- Three years ago, 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”

© 2009 Sirius Computer Solutions 11

---

---

---


---

---

---

---

---

**First Year’s Linux DR Test...** 

- First DR test with z/VM and Linux
- Fall ‘07
  - Monday to Thursday
- Customer personnel
  - NOT onsite at DR vendor site
  - NOT at their normal work site
  - Across town with only DR doc
  - Remote access to DR vendor & systems

© 2009 Sirius Computer Solutions 12

---

---

---

---

---

---

---

---

**First Year's Linux DR Test...**

- 2 z/OS LPARs
- 1 z/VM LPAR
- 20+ Linux Guests
  - 4-8 Production
  - "Why not bring them all up"
- Still many "Distributed Systems"
  - Several moved to Linux on Z
  - Fewer unique environments

© 2009 Sirius Computer Solutions 13

---

---

---

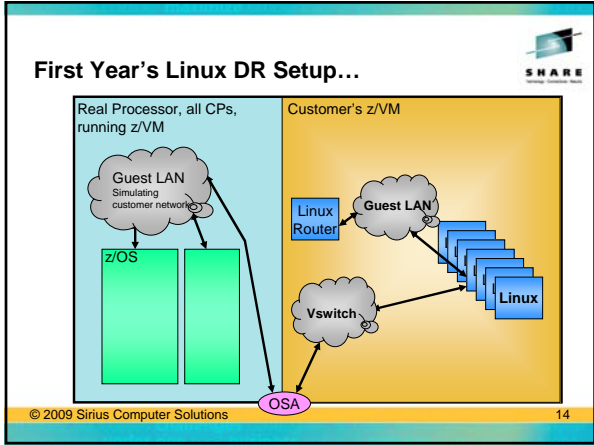
---

---

---

---

---




---

---

---

---

---

---

---

---

**First Year's Linux DR 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"

© 2009 Sirius Computer Solutions 15

---

---

---

---



---

---

---

---

### First Problem

- Monday late afternoon
- Performance!
- Customer's VM image only given less than 1/2 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

© 2009 Sirius Computer Solutions 16

---

---

---

---

---

---



---

---

---

---

### 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's VM!
- Competing with z/OS for cycles
- Can't shutdown z/OS systems to reconfigure LPARs

© 2009 Sirius Computer Solutions 17

---

---

---

---

---

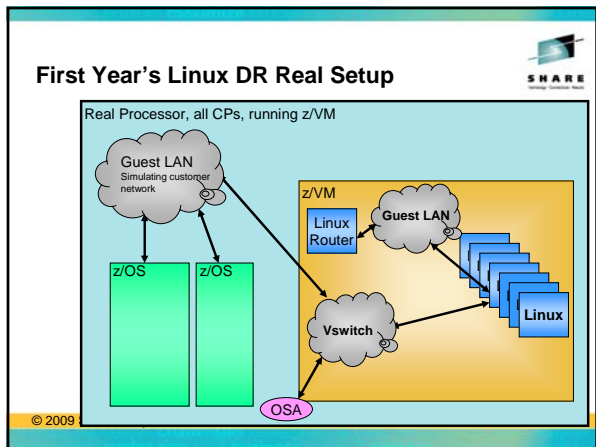
---

---

---

---

---




---

---

---

---

---

---


---

---

---

---

**Second Problem – Why?**



- Customer VM & CMS runs great, Linux runs terrible
- Plenty of real memory, 2 CPs running 100%
- Why? Hardware vs. Software simulation

© 2009 Sirius Computer Solutions 19

---

---

---

---


---

---

---

---

**Second Problem – Why?**



- System Z hardware only supports 2 levels of virtualization – 2 levels of SIE
  - LPAR support uses the 1<sup>st</sup> level SIE to run the LPARs (DR Vendor's VM system)
  - The DR vendor's VM system (1<sup>st</sup> level) uses the 2<sup>nd</sup> level SIE to run the 2<sup>nd</sup> level systems (the customer's z/OS & VM images)
- The customer's VM system (2<sup>nd</sup> level) cannot use SIE to manage the Linux guests (3<sup>rd</sup> level)
- CP must simulate all privops and management tasks
- The Fix? Avoid that 2<sup>nd</sup> level of VM

© 2009 Sirius Computer Solutions 20

---

---

---

---


---

---

---

---

**Second Problem – The Fix**



- Option #1 - Put your VM in an LPAR on the DR vendor's machine
  - NOT under the vendor's VM!
  - 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

© 2009 Sirius Computer Solutions 21

---

---

---

---

---

---

---

---

## Second Problem – The Fix



- **Option #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 testing purposes

---

---

---

---

---

---

---

---

## 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

---

---

---

---

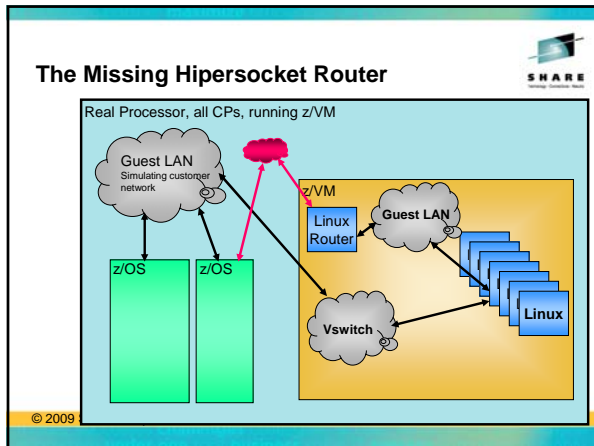
---

---

---

---






---

---

---

---

---

---

---

---

---

---

- ### Second 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 "critical" production machines
    - Distributed Systems – "Working on it"
  - **One "small" problem**
    - Missing Vswitch controller userids
- © 2009 Sirius Computer Solutions 26

---

---

---

---

---

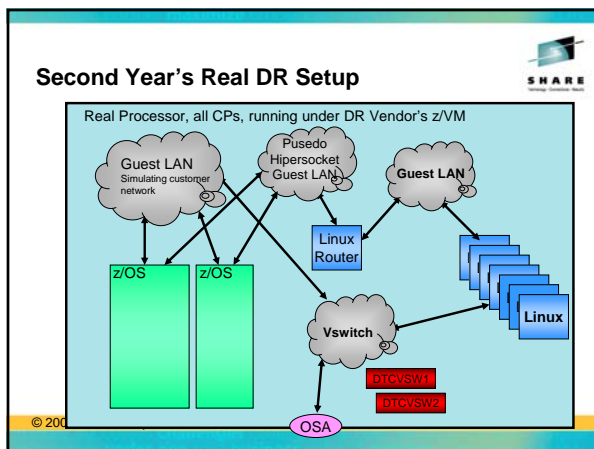
---

---

---

---

---




---

---

---

---

---

---


---

---

---

---

**Third Year's DR Test**



- **Customer's VM in a DR LPAR**
  - Cost a little more, but...
  - Only changes were OSA addresses
- **Bring up all the Linux machines**
  - 20+ Production, Dev, Q/A and Test

© 2009 Sirius Computer Solutions 28

---

---

---

---



---

---

---

---

**Third Year's DR Test**

- **Monday morning:**
  - Start restoring tapes
- **Monday early afternoon:**
  - 2 z/OS LPARs – Up and running
  - 1 z/VM LPAR – Up and running
  - All 20+ Linux machines – Up and running
  - Distributed Systems – “Working on it”
- **Two IPLs of VM**
  - First to change the OSA & Hipersocket addresses
  - Second to bring everything up

© 2009 Sirius Computer Solutions 29

---

---

---

---


---

---

---

---

**Third Year's DR Test Results**



- **Customer very pleased with DR speed**
  - Customer VM in an LPAR cost a little more, but...
  - Well worth the investment
- **Evaluate migrating remaining non-mainframe workload to Linux on z**
  - DR recovery speed is only one benefit
  - A lot of territorial resistance – FUD  
Fear – Uncertainty – Doubt
- **DR the Development and Test environments also**

© 2009 Sirius Computer Solutions 30

---

---

---


---

---

---

---

---

**Lessons Learned** 

- 1. Communicate!**
  - What we thought was going to be the configuration, wasn't, more than once!
  - 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
  - Fight for what you need, not just what's easy or cheap

© 2009 Sirius Computer Solutions 31

---

---

---

---

---


---

---

---

---

---

**Lessons Learned** 

- 3. Be SURE of your DR configuration**
  - For each LPAR, for each VM, for each Linux
  - 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 3<sup>rd</sup> level Linux**
  - #1 issue: No Linux under VM under VM
  - Your options are:
    1. Run your VM in a DR LPAR
    - or -
    2. Run your Linux directly under the DR VM

© 2009 Sirius Computer Solutions 32

---

---

---

---

---


---

---

---

---

---

**Lessons Learned** 

- 5. Check your configuration when you get on the DR machine**
  - Hardware, memory, network
  - Don't assume you have all you asked for
  - Don't assume it's set up right
  - Don't assume you have what you did last time
  - 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

© 2009 Sirius Computer Solutions 33

---

---

---

---

---

---

---

---

---

---

All Too Easy To Do...



**Our Disaster Recovery Plan Goes Something Like This...**



© 2009 Sirius Computer Solutions 34

---

---

---

---

---

---

---

---



**z/VM and Linux  
Disaster Recovery –  
A Customer Experience**

Lee Stewart  
Sirius Computer Solutions (DSP)  
lee.stewart@siriuscom.com

Session 9210



---

---

---

---

---

---

---

---