#SHAREorg

SHARE
Technology · Connections · Results

Disaster Recovery of Linux on System Z

Deric Abel
America First Federal Credit Union

Wednesday, February 6th, 2013
Session 12147

http://www.linkedin.com/pub/deric-abel/1/738/581/

SHARE in San Francisco 2013

Agenda

- Introductions
- Mainframe at AFCU
- DR environment
- Automation
  - Scripts and configuration files
- Testing

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Speaker Introduction

- Deric Abel
- I've been in IT since 1997
- First installed Linux as a High School project my senior year (1999)
- Hired as a Linux Admin in 2000
- First experience with Virtualization in 2005
- Hired at America First Credit Union as a z/Linux admin in 2008
- Joined the zNextgen group and attended my first SHARE conference in 2008
- Currently serving as a Deputy Project Manager for zNextgen

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Disclaimer

- This presentation is for information purposes only.
- This is not an endorsement of America First Credit Union.
- Every environment is different and the scripts used in this presentation most likely will not work in your environment.

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Current Environment

- Two z10
  - Z02 at Primary Data Center
  - O01 at Backup Data Center
    - CBU to Z02 during DR test or actual DR event.
- Primary z10
  - 4 z/VM LPARs
    - 2 Production
    - 2 Test/Development
  - 4 IFLs and 2 CPs
  - 56GB Memory
  - Level 9 facility
    - Full redundancy in UPS, switch, and generator systems
    - Building is “base isolated” to withstand a horizontal shift

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Primary Mainframe Environment

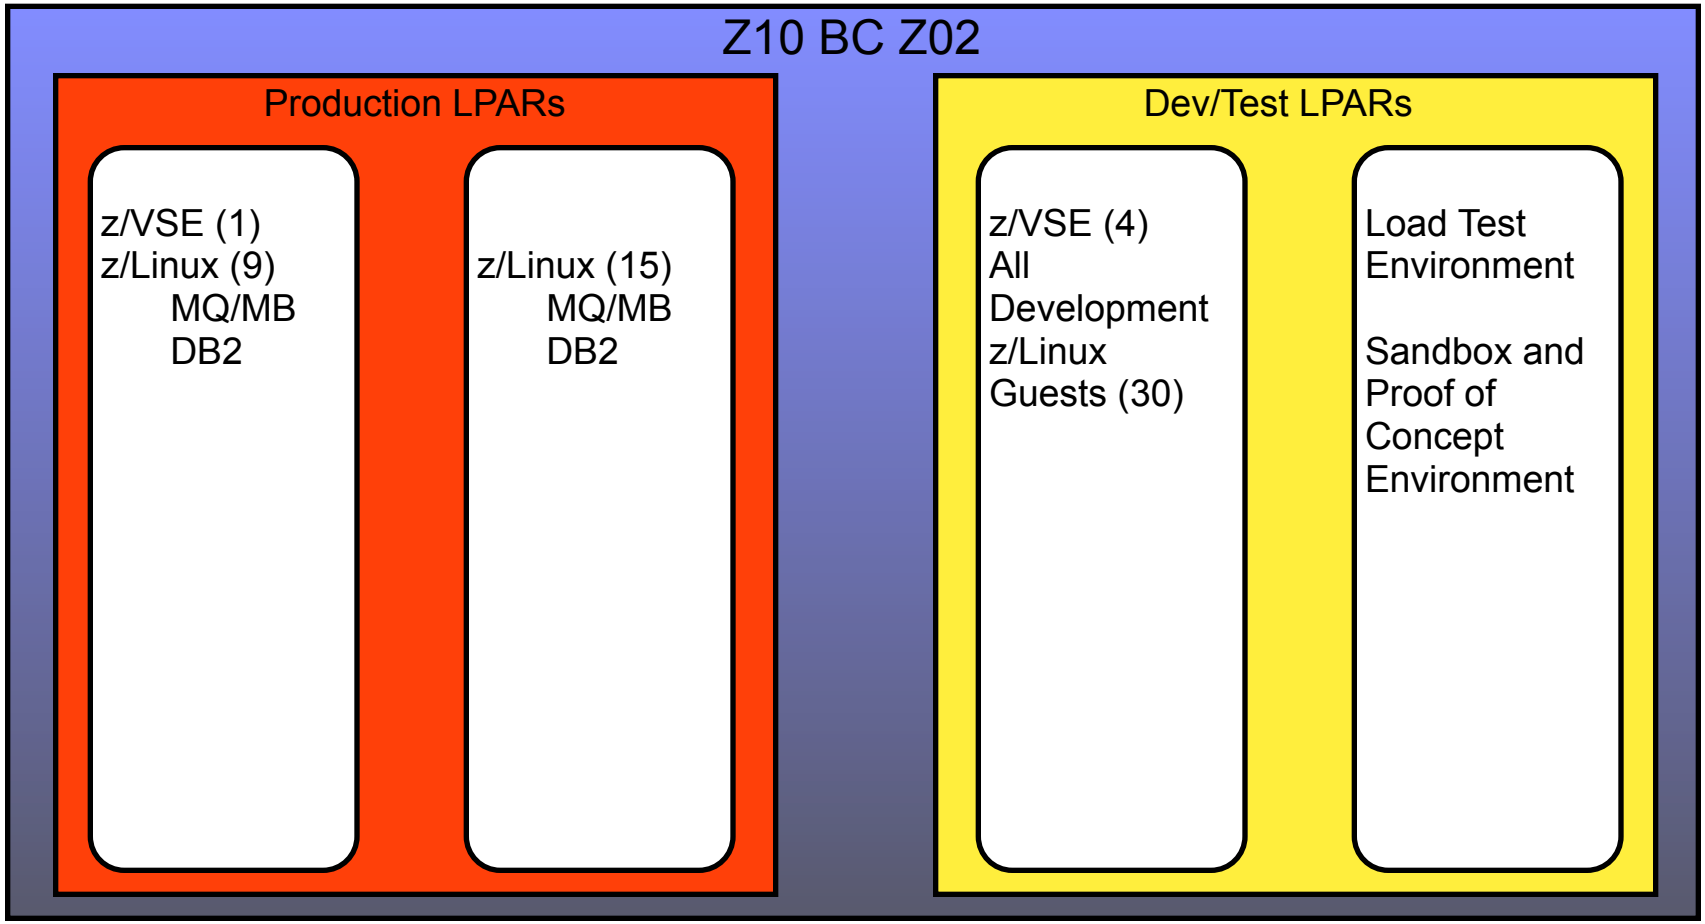

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Primary Mainframe Environment

- Linux OS installed on ECKD disk
- Middleware and Database volumes use SAN disk over FCP
- Hipersocket network used for guest to guest communication
- z/VM LDAP server for Linux users and groups
- DNS servers for mainframe reside on Linux servers (these point back to corporate name servers)

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

DR Setup

- Replication between primary and backup data centers us Global Mirror Replication (TPC-R)
- Backup mainframe
  - LPAR names are different between data centers
    - Primary DC LPAR name starts with DC
    - Backup DC LPAR name starts with SLC
      - *Guests see the LPAR name at boot up and loads its config based on name of LPAR*
    - SAN target WWPN and LUN differ between primary and backup data centers
      - *Guests automatically configure WWPN and LUN based on LPAR it comes up in*
  - Hipersocket address DO NOT change at backup data center
  - External network address DO change at backup data center

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

DR Setup

- From an OS perspective everything is automatic
- Manual processes
  - Update DNS to reflect DR IP addresses
  - Update middleware to point to external(to mainframe) databases

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

DR Setup

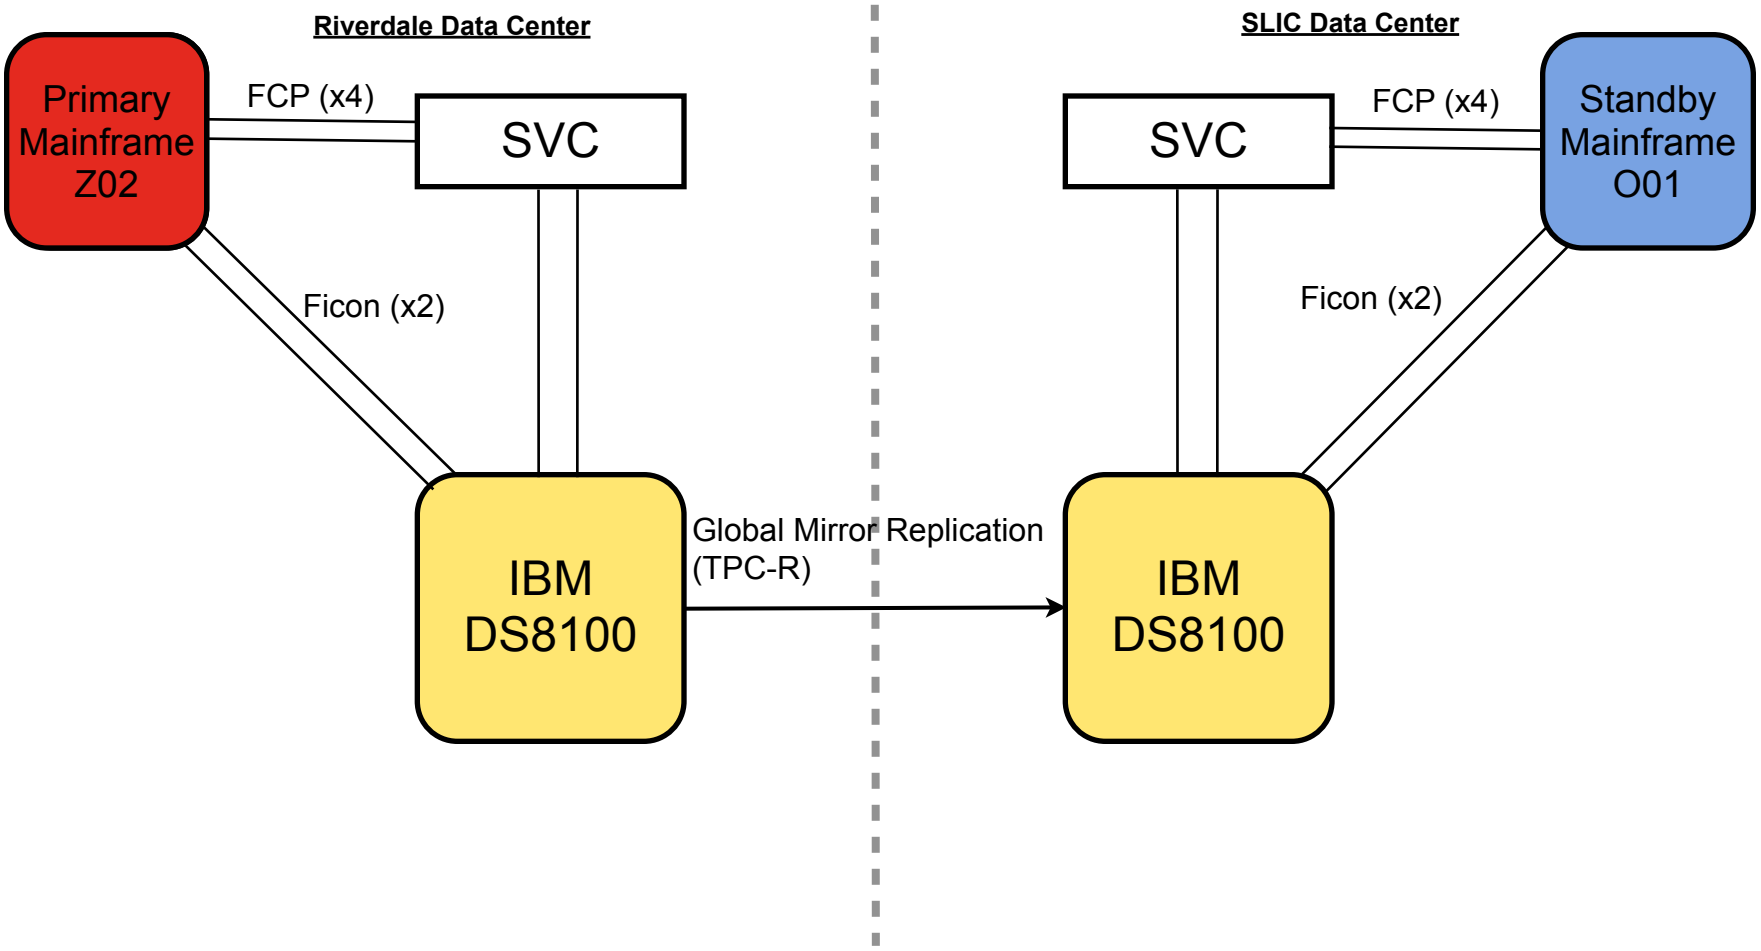

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

DR Event

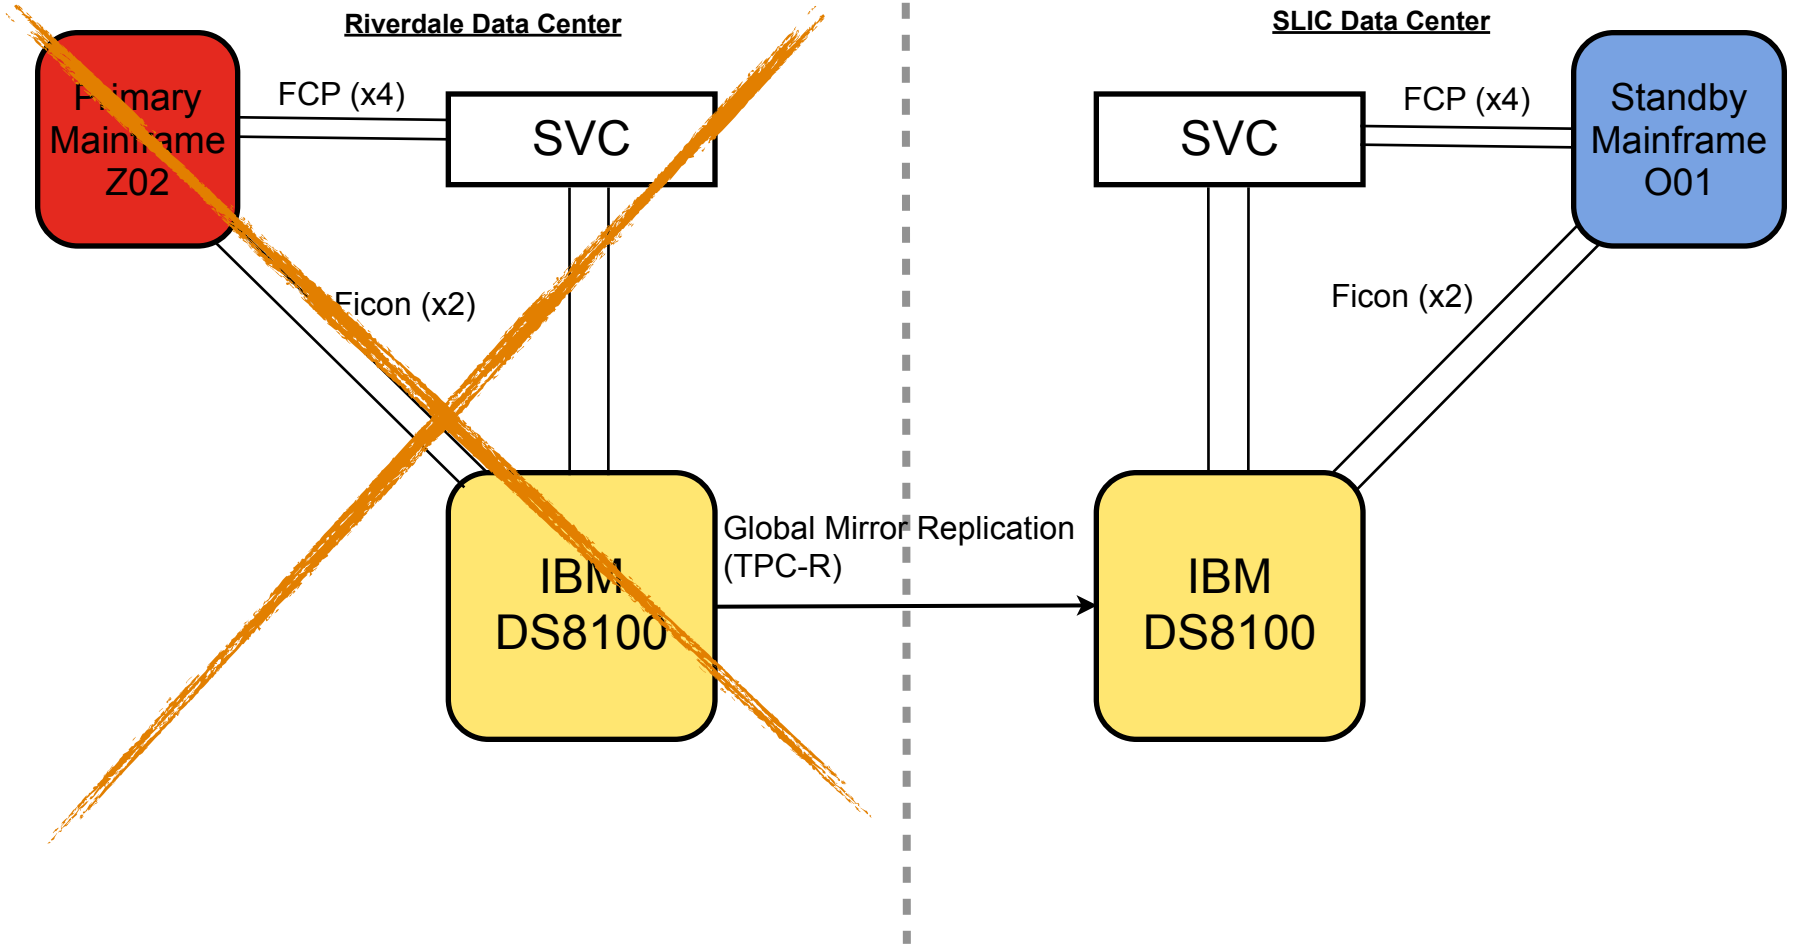

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

DR Event

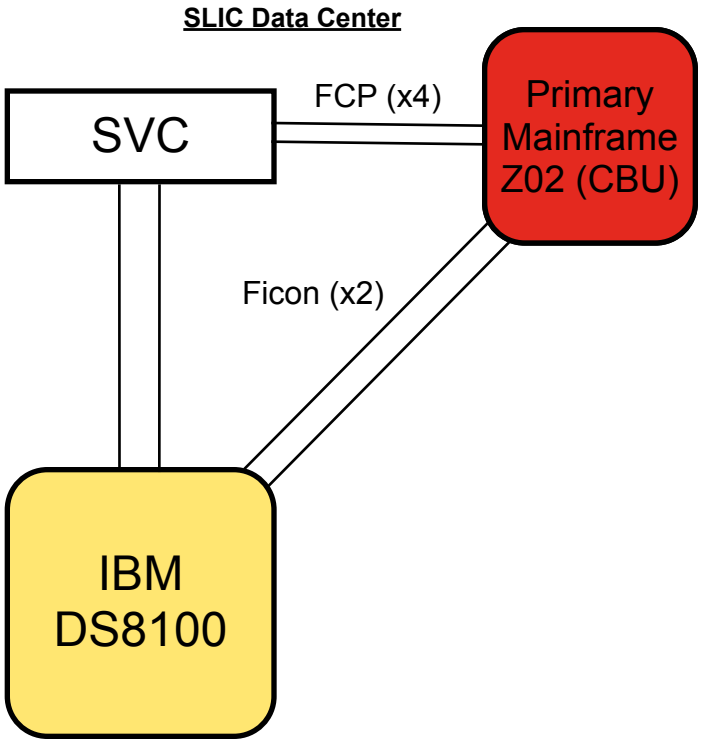

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation *(From Nationwide Insurance DR presentation)

- Avoid manual processes
  - Dependence on key individuals
  - Prone to mistakes
  - Slow
- Automated processes
  - Requires only basic knowledge of environment and technologies in use
  - Accuracy
  - Repeatable
  - Faster
  - Does not mean build it once then ignore; Requires regular review and updates

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation

- Linux configuration happens at boot
  - As part of boot, the script boot.config is run
    - Identifies LPAR by interacting with CP using VMCP
    - Configures WWPN and LUN for fcp volumes
    - Configures network IP information (physical and hipersocket)
    - Linux PARM file is stored on a commonly accessible CMS disk
      - *Using cmsfscat, the [guestid].parm is written to /tmp/sourceinfo*

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation

LNXLP001.PARM

```
HOST=lnxlp001
DCIP=10.215.121.17
DCRT=10.215.121.1
DCHSI=10.215.9.110
SLCIP=10.225.121.17
SLCRT=10.225.121.1
SLCHSI=10.215.9.110
DCFCP='2000 2100'
DCWWPN_2000="5005076801302dc0
5005076801202dc0
5005076801302da6
5005076801202da6"
DCWWPN_2100="5005076801402dc0
5005076801102dc0
5005076801402da6
5005076801102da6"
DCLUNS="0000 0001"
SLCFCP="2000 2100"
SLCWWPN_2000="5005076801203a48
5005076801303a48
5005076801203a55
5005076801303a55"
SLCWWPN_2100="5005076801103a48
5005076801403a48
5005076801103a55
5005076801403a55"
SLCLUNS="0000 0001"
ENV=PROD
```


Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation
boot.config

```
#! /bin/sh
#
### BEGIN INIT INFO
# Provides:           boot.config
# Required-Start:     boot.udev
# Required-Stop:
# Should-Start:
# X-Start-Before:      boot.multipath boot.lvm
# Default-Start:      B
# Default-Stop:
# Description:        install config files
### END INIT INFO

. /etc/rc.status
```

```
build_nic_config () {

  echo "BOOTPROTO='static'" > $1
  echo "UNIQUE=''" >> $1
  echo "STARTMODE='auto'" >> $1
  echo "IPADDR='$2'" >> $1
if [ -z $3 ] ; then
  echo "NETMASK='255.255.255.0'" >> $1
else
  echo "NETMASK='$3'" >> $1
fi
  echo "NETWORK=''" >> $1
  echo "BROADCAST=''" >> $1
  echo "ETHTOOL_OPTIONS=''" >> $1
  echo "MTU=''" >> $1
  echo "NAME=''" >> $1
  echo "REMOTE_IPADDR=''" >> $1
  echo "USERCONTROL='no'" >> $1
  echo "PREFIXLEN=''" >> $1
}
```


Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation
boot.config (cont.)

```
case "$1" in
  start)
# modprobe required just in case
modprobe vmcp
sleep 1
/sbin/vmcp "link * 191 191 rr"
sleep 2
echo "1" > /sys/bus/ccw/devices/0.0.0191/online
sleep 2
PARMDEV=`grep 191 /proc/dasd/devices|awk '{print $7}'`
QUSERID=`/sbin/vmcp query userid`
GUEST=`echo $QUSERID|awk '{print $1}'`
LPAR=`echo $QUSERID|awk '{print $3}'`
echo "GUEST=$GUEST" > /tmp/sourceinfo
echo "LPAR=$LPAR" >> /tmp/sourceinfo
cmsfscat -d /dev/$PARMDEV -a ${GUEST}.parm >> /tmp/sourceinfo
echo "0" > /sys/bus/ccw/devices/0.0.0191/online
/sbin/vmcp "det 191"

. /tmp/sourceinfo
```


Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation
boot.config (cont.)

```
case "$ENV" in
        PROD)
                CLR="41"; #Red
                ;;
        TEST)
                CLR="42"; #Green
                ;;
        DR)
                CLR="46"; #Turq
                ;;
esac
echo "CLR=$CLR" >> /tmp/sourceinfo
case "$LPAR" in
        DC*)
                SITE="DC";
                ln -sf /etc/hosts.dc /etc/hosts
                ;;
        SLC*)
                echo DR="[DR]" >> /tmp/sourceinfo;
                ln -sf /etc/hosts.slc /etc/hosts
                SITE="SLC";
                ;;
esac
```


Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation
boot.config (cont.)

```
IP=${SITE}IP
MASK=${SITE}MASK
HIP=${SITE}HSI
RT=${SITE}RT
FCP=${SITE}FCP
WWPN=${SITE}WWPN
LUNS=${SITE}LUNS

for _fcp in ${!FCP}; do
  zfcp_host_configure 0.0.${_fcp} 1
  for _lun in ${!LUNS}; do
      _port_list=${WWPN}_${_fcp}
      for _port in ${!_port_list}; do
          zfcp_disk_configure 0.0.${_fcp} 0x${_port} 0x${_lun}000000000000 1
      done
  done
done

#Build IP config files

build_nic_config "/etc/sysconfig/network/ifcfg-qeth-bus-ccw-0.0.1000" ${!IP} ${!MASK}
build_nic_config "/etc/sysconfig/network/ifcfg-hsi-bus-ccw-0.0.e000" ${!HIP}
```


Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation
boot.config (cont.)

```
echo "default ${!RT} - -" > /etc/sysconfig/network/routes
echo "$HOST.systems.americafirst.com" > /etc/HOSTNAME
hostname $HOST
#Build MOTD
echo "zLinux  $LPAR  $GUEST  `hostname`" > /etc/motd
echo "`head -1 /etc/SuSE-release`" >> /etc/motd
echo "`uname -rv`" >> /etc/motd

[ -f /etc/motd.skel ] && cat /etc/motd.skel >> /etc/motd
rc_status -v
;;
  stop)
    #do nothing
    rc_status -v
    ;;
  *)
   echo "Usage: $0 {start|stop}"
   exit 1
   ;;
esac

rc_exit
```


Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation sourceinfo

```
GUEST=LNXLP001
LPAR=DCPROD1
HOST=lnxlp001
DCIP=10.215.121.17
DCRT=10.215.121.1
DCHSI=10.215.9.110
SLCIP=10.225.121.17
SLCRT=10.225.121.1
SLCHSI=10.215.9.110
DCFCP='2000 2100'
DCWWPN_2000="5005076801302dc0
5005076801202dc0
5005076801302da6
5005076801202da6"
DCWWPN_2100="5005076801402dc0
5005076801102dc0
5005076801402da6
5005076801102da6"
DCLUNS="0000 0001"
SLCFCP="2000 2100"
SLCWWPN_2000="5005076801203a48
5005076801303a48
5005076801203a55
5005076801303a55"
SLCWWPN_2100="5005076801103a48
5005076801403a48
5005076801103a55
5005076801403a55"
SLCLUNS="0000 0001"
ENV=PROD
CLR=41
```


Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation
MOTD and Prompt

Live Production (non-DR)

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation
MOTD and Prompt

Guest in DR Mode

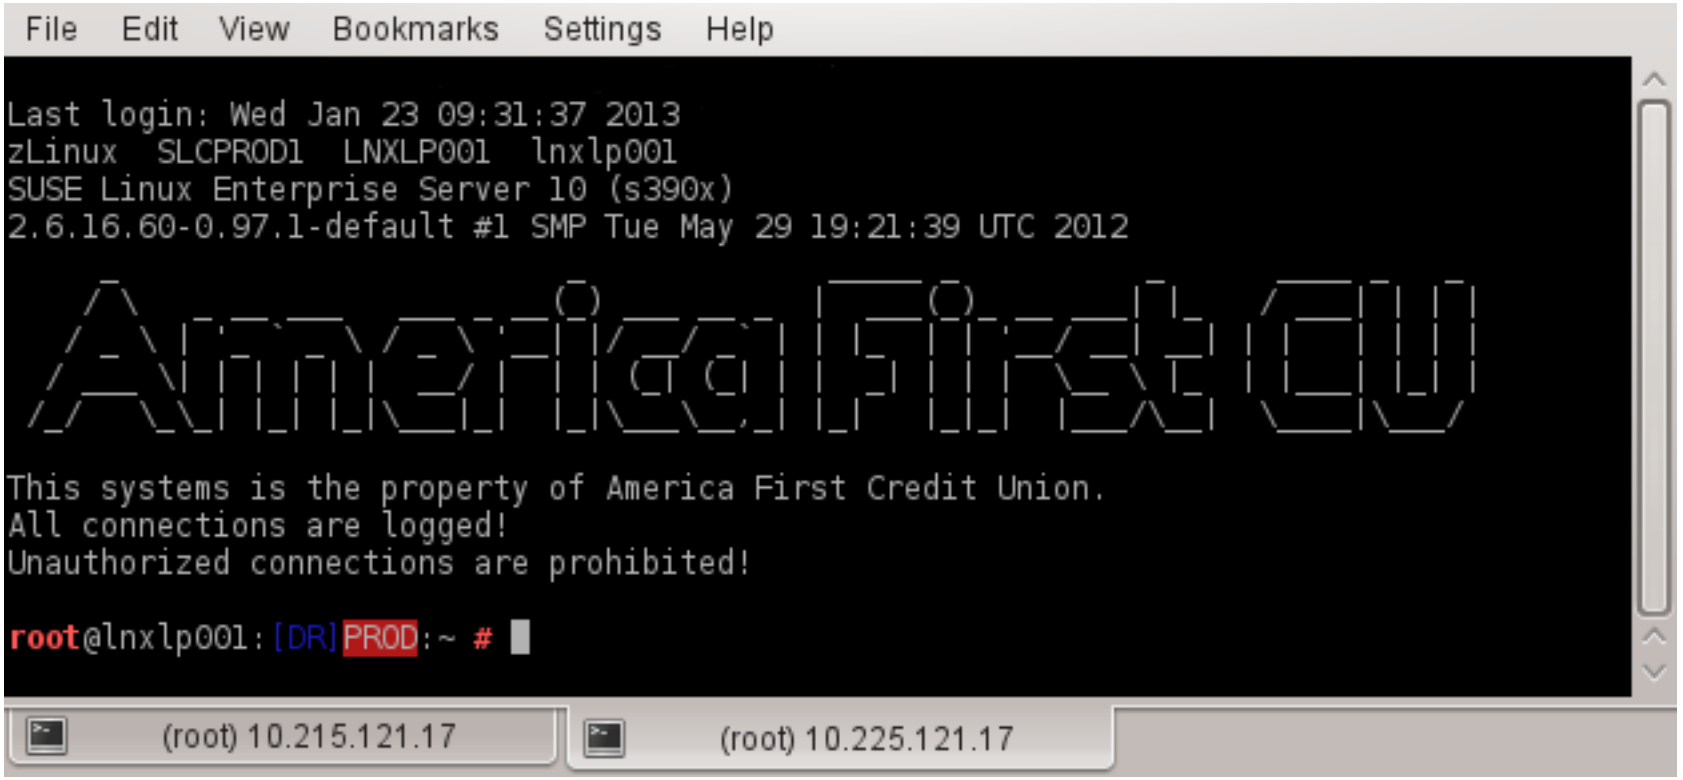

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Automation
bash.bashrc.local

```
source /tmp/sourceinfo
_t=""
_e="\[\e[${CLR}m\]$ENV\[\e[0m\]"
_dr="\[\e[34m\]$DR\[\e[0m\]"
if test "$UID" = 0 ; then
   _u="\[\e[1;31m\]\u\[\e[0m\]@\h"
   _p="\[\e[1;31m\] #\[\e[0m\]"
else
   _u="\u@\h"
   _p=">"
   if test \( "$TERM" = "xterm" -o "${TERM#screen}" != "$TERM" \) \
      -a -z "$EMACS" -a -z "$MC_SID" -a -n "$DISPLAY"
   then
      _t="\$(ppwd \l)"
   fi
fi
# With full path on prompt
PS1="${_t}${_u}:${_dr}${_e}:\w${_p} "
```


Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Testing

- Test annually
- Document everything during your test.
  - Track all issues found during test.
    - Included issues that were resolved.

- In our DR test we take a snapshot (Point-in-Time Copy) of our production data and use that as our DR test dataset.
- We test our DR environment in a “test” mode only.
  - Production traffic continues to our non-DR mainframe.
- Once test is complete, DR test dataset is discarded.

Complete your sessions evaluation online at SHARE.org/SanFranciscoEval

Questions

#SHAREorg

SHARE
Technology · Connections · Results

Disaster Recovery of Linux on System Z

Deric Abel
America First Federal Credit Union

Wednesday, February 6th, 2013
Session 12147

http://www.linkedin.com/pub/deric-abel/1/738/581/

SHARE in San Francisco 2013