BEST PRACTICES

FOR RED HAT ENTERPRISE LINUX

ON SYSTEM Z

Brad Hinson

bhinson@redhat.com

IBM – RED HAT

RHEL SOE AUDIT REPORT

Project Delivery Documentation

Red Hat Asia Pacifc
Level 5

455 Bourke Street

Melbourne VIC 3000

Trademarks

Trademarked names may appear throughout this document.

Rather than list the names and entities that own the trademarks or insert a trademark symbol with each mention of the trademarked name, the names are used only for editorial purposes and to the benefit of the trademark owner with no intention of infringing upon that trademark.

TABLE OF CONTENTS

1. Executive Summary

1.1. Background

IBM has requested Red Hat Consulting Services to undertake a review of the current Red Hat Enterprise Linux environment. The environment that was nominated for review was the Acme Application environment that is currently being migrated from Intel x86 based servers to the mainframe, z-Series. The review is required to determine if the deployed environments are appropriate for IBM/Acme's requirements. This report outlines the findings of that work.

1.2. Project Goal

The goal of the project is to review the Red Hat Enterprise Linux environment currently in place at ACME Application, with a view to highlight any concerns and make recommendations for remediation of those concerns.

1.3. Key Assumptions

Red Hat assumes that the information provided by Acme and the outsourcing partner organisation IBM are true and accurate for their production environment.

1.4. High level Findings

- z/VM requires fine tuning
- Virtual Machine's sizing needs to be reviewed
- RHEL requires tunings
- Puppet configuration requires optimisations

2. Key Recommendations

3. Findings

These findings identify areas of the installed environments that do not meet best (should this be leading practice instead of best practice?) practice. It is understood that an environment may deviate from generic best practice due the specialised requirements of the solution.

Each item will be broken down into the following areas:

Details: Outlines the nature of the issue
Red Hat's Recommendation: Red Hat's suggestion to resolve the issue
Probability: The likelihood that this item will cause a problem
Impact: The impact that the item will have to the operations or security of the server.
Affected Environments: Environments that were affected by the finding.

3.1. z/VM Configuration

a) Startup and Shutdown

1. Startup script

Details: There is no startup script for setting the initial action when a virtual machine guest logs on to z/VM.

Red Hat's Recommendation: Create a PROFILE EXEC file on the LINUXADM user's 191/a disk, which is automatically utilized by each logged in user. This should contain the following lines:

```
/* Profile Exec */
'CP SET PF11 RETRIEVE FORWARD'
'CP SET PF12 RETRIEVE BACKWARD'
'CP SET RUN ON'
```

Note that the first line must be enclosed in /* */ C-style comments. The 'CP SET RUN ON' allows the virtual machine to continue running automatically when logging into a disconnected virtual machine.

Probability: High. Affects each login to a virtual machine.

Impact: High. Without this startup file, a virtual machine could be paused unintentionally, causing a temporary outage with each 3270 login.

2. Virtual machine startup menu

Details: The startup script discussed above can also be used to present the user with a common menu of actions at login.

Red Hat's Recommendation: In the shared PROFILE EXEC mentioned above, add a menu written in REXX. A detailed sample is included in the Appendix 1.

Probability: Medium. This can lower the learning curve for system administrators not familiar with the 3270 console. Additionally, with this in place, z/VM can now start the virtual machine automatically at startup without manual interaction, reducing downtime after an outage.

Impact: Medium. Does not affect system performance, but increases usability.

3. z/VM Shutdown

Details: z/VM can signal a virtual machine to shutdown, then wait for it. This is useful when there are many systems to halt, for example when z/VM is shutdown for maintenance. The current setting is for z/VM to wait 30 seconds after signaling a virtual machine shutdown before considering the guest down.

Red Hat's Recommendation: Increase this setting to at least 180 seconds (but no more than 300), to allow for middleware to shutdown gracefully. The following line should be added to the user AUTOLOG1's PROFILE EXEC startup file:

```
'CP SET SIGNAL SHUTDOWN 180'
```

Probability: Low. Systems are not shutdown frequently.

Impact: Low. Systems can be alternatively be shutdown manually at the console, remotely over SSH, or through the Satellite.

b) z/VM Memory Tuning

1. Set memory and scheduling tunables in system configuration

Details: z/VM should be configured for optimal memory usage based on the proposed system workload. This involves tuning the minidisk cache, storage pool buffer sizes, and VDISK limits.

Red Hat's Recommendation: Configure minidisk cache so that it's taken from central storage instead of expanded, then limit the maximum amount. Add the following lines to the user AUTOLOG1's PROFILE EXEC:

```
'CP SET MDC XSTOR 0M 0M'
'CP SET MDC STORE 0M 256M'
```

For more information on minidisk cache recommendations, see
http://www.vm.ibm.com/perf/tips/prgmdcar.html

Next, configure z/VM resource management and scheduling values based on recommendations for Linux virtual machines. Add the following line to the user AUTOLOG1's PROFILE EXEC:

```
'CP SET SRM STORBUF 300 250 200'
```

Probability: High. Affects scheduling and memory management for all virtual machine under z/VM.

Impact: High. Setting these values ensures Linux is scheduled properly under z/VM, and minidisk cache usage is optimized.

2. Add additional volumes for paging and spool

Details: z/VM is preconfigured with only a single DASD volume for paging, and a single DASD volume for spool space. Based on the projected workload, this will not sustain the environment.

Red Hat's Recommendation: Format additional DASD to be used as paging volumes, then add their labels to the "CP_Owned" section of SYSTEM CONFIG. Paging volumes should be equal in geometry and number of cylinders. To calculate the total number of volumes needed, add the total projected memory size for all guests, then subtract the memory size of the LPAR. Then double this number, because the z/VM paging subsystem performs optimally when its volumes are less than 50% used.

Adding an additional volume for spool is optional but recommended. The current observed spool usage was high, which could prevent manual installations; however, as kickstart is used, this becomes less important. Consider adding spool if usage is above 90%.

Probability: Paging: High. As Linux virtual machines are added, their memory becomes eligible for paging. z/VM must have the capacity to page these to disk. Spool: Low. Usage should be monitored and additional volumes added above 90%.

Impact: High. Affects system response time, and allows z/VM to overcommit memory.

3. Increase VDISK limits

Details: The default values for VDISK system and user limits may not satisfy the total VDISK required by the Linux guests.

Red Hat's Recommendation: Increase the VDISK limits. In the "Features" section of SYSTEM CONFIG, make the following change:

```
Vdisk ,
   Userlim infinite ,      /* Maximum vdisk allowed per user */
   Syslim  infinite        /* Maximum vdisk allowed for system */
```

Probability: High. Every Linux guest should be configured to have VDISK, so default limit is reached too quickly.

Impact: High. Allows better swap and memory management for Linux guests.

c) DASD Disk Configuration

1. Relabel system disks

Details: z/VM uses a default disk label for all system disks. This could conflict with another z/VM instance on the same storage array, especially at startup.

Red Hat's Recommendation: Relabel the five system disks (610RES, 610SPL, 610PAG, 610W01, and 610W02, if they exist) to something unique to each environment (i.e. B02RES, B02SPL, etc). This procedure is outlined in the following PDF:

http://www.redbooks.ibm.com/redbooks/pdfs/sg247932.pdf

(see section 4.11 "Relabeling system volumes," page 66, PDF page 86)

Probability: Low. Issue can be avoided if IOCDS definition prevents one LPAR from seeing disks used in another LPAR.

Impact: Medium. Only applicable at z/VM boot time, but if not properly configured, could prevent system from starting.

2. Use non-sequential DASD

Details: Since DASD is emulated as partial storage on large real disks across multiple control units, there is the potential for channel bottleneck if sequential disks are given to a single system, including z/VM and Linux.

Red Hat's Recommendation: Check with storage array vendor on whether SAN is implemented with internal striping. If not, any DASD used for Linux, or z/VM paging, should be striped across control units instead of allocated sequentially for any one purpose.

Probability: Low. Most recent EMC and IBM storage arrays implement internal striping.

Impact: High. Severely impacts performance as multiple DASD are contained onto a single physical disk.

d) Networking

Backup devices for virtual switch

Details: The VSWMGMT1 virtual switch has only one OSA real device (rdev). It is not protected against external switch failure.

Red Hat's Recommendation: Add a second rdev to this VSWITCH.

Probability: Low. As discussed, this is currently being planned.

Impact: Medium. Without multiple rdev statements, the VSWITCH is not protected against network failure.

e) Security

Configure hardware cryptographic assist

Details: With hardware crypto, secure transaction processing is offloaded instead of handled by the IFL. This can be a big performance benefit. Crypto is configured for the LPAR, but is not configured for each virtual machine.

Red Hat's Recommendation: Configure each guest to access the Crypto device. Add the following line to the virtual machine guest definition, preferably into a shared profile:

```
'CRYPT APVIRT'
```

Probability: High. Initial query from guest indicated Crypto was not yet defined, while query from user MAINT showed Crypto device is available.

Impact: Medium. Depending on the amount of SSL/RSA/DSA secure processing, this could free CPU cycles from the IFL, giving a higher consolidation ratio of virtual machines per LPAR.

3.2. Virtual Machine Sizing

a) Add Second Virtual CPU

Details: The LPAR has four IFLs allocated, but each individual guest has only one virtual CPU defined. This limits the amount of multiprocessing for multi-threaded applications, like WebSphere.

Red Hat's Recommendation: Define at least two virtual CPUs to each guest, but no more than the number of IFLs for this LPAR (four). At this time the recommended number of virtual CPUs is two.

Probability: Low. Only development environment was observed. Production environment may already have two virtual CPUs defined.

Impact: High. WebSphere has a significant benefit from multiple CPUs in Linux.

b) SWAPGEN Script

Details: This script, developed by Sine Nomine Associates, is used to prepare a VDISK for use as a swap device in Linux (http://www.sinenomine.net/products/vm/swapgen). It is called from the virtual machine's PROFILE EXEC so the VDISK is created before the system boots.

Red Hat's Recommendation: Download the SWAPGEN script, then make it available on the user LINUXADM's 191/a disk. Use SWAPGEN to create two VDISK swap spaces, at addresses 300 and 301. The first VDISK is 256MB, and the second is 512MB. Add the following lines to the user LINUXADM's PROFILE EXEC:

```
'SWAPGEN 300 524288'
'SWAPGEN 301 1048576'
```

Include 300-301 in the DASD list (in the CMS CONF FILE before installation, or /etc/modprobe.conf afterwards). In Linux, use pri= in the /etc/fstab options to give the smallest VDISK swap the highest priority, followed by the larger VDISK, then swap space on DASD with the lowest priority.

Probability: High. Initial audit showed only DASD swap was being used.

Impact: High. Using VDISK swap can significantly improve Linux swap performance since VDISK is implemented in real memory.

c) Disable Minidisk Cache for Swap Device

Details: Minidisk swap is enabled for all disks unless explicitly disabled. Since swap devices are designed for many writes and few reads, minidisk cache should not be used.

Red Hat's Recommendation: Instead of using an LVM partition for swap, split the existing minidisk into two disks, with one dedicated as a swap device. In the user directory, add the following line after the swap minidisk:

```
MINIOPT NOMDC
```

Probability: Low. Swap should happen on VDISK first, then DASD based on priority. This is only for very high memory pressure situations.

Impact: Medium. Depending on swap usage, this could have a positive impact on Linux memory subsystem performance.

3.3. *Red Hat Enterprise Linux Configuration*

a) Pre-emptive Support and Troubleshooting

1. DASD dump

Details: In cases where Red Hat/IBM support requires a kernel core image for diagnostic troubleshooting, for example after a system crash or performance-related problem, DASD can be used to obtain a vmcore dump of the virtual machine memory. This requires a minidisk at least as large as the virtual machine memory size.

In cases where disk space is a shortage, if every virtual machine cannot have a dedicated minidisk for DASD dump, a shared minidisk can be used for multiple guests. However, if any two guests experience a system crash at the same time, both vmcore images will be unusable. This should only be used when a single vmcore image is captured at a time.

Red Hat's Recommendation: Configure DASD dump as outlined in the IBM documentation (see “Using the Dump Tools” http://public.dhe.ibm.com/software/dw/linux390/docu/l26cdt02.pdf)

Probability: Low. System crash and kernel vmcore capture is a rare occurrence.

Impact: High. Providing vmcore data to Red Hat/IBM support greatly expedites issue resolution.

2. IUCV terminal

Details: System z uses an internal network called IUCV to implement a TTY terminal client and server. This terminal can replace the 3270 console for troubleshooting a virtual machine when TCP/IP and SSH are not available, for example during system startup. This terminal also allows traditional Linux text editors to run, such as **vi** and **emacs**, instead of relying on line-mode editors in the 3270 console. See Appendix 2 for graphical depiction.

Red Hat's Recommendation: Configure the IUCV terminal as outlined in the IBM documentation (See “How to Set up a Terminal Server Environment” http://public.dhe.ibm.com/software/dw/linux390/docu/l26dht00.pdf)

Probability: Low. This is only used when troubleshooting a problem.

Impact: Medium. Replaces the need to train system administrators on 3270 console usage.

3. Rescue mode

Details: In the event the virtual machine cannot startup normally, Red Hat Enterprise Linux can be configured to start in rescue mode, which is a failsafe mode used for system recovery.

Red Hat's Recommendation: Configure system startup with the option to load rescue mode, to be used if the system cannot boot normally. This procedure is described in the following document:

http://www.redbooks.ibm.com/redbooks/pdfs/sg247932.pdf

(see section 13.5 “Rescuing a Linux system,” page 219, PDF page 239).

In addition, this can be added to a common startup menu as an additional boot option. See section “Virtual machine startup menu” under *z/VM Configuration, Startup and shutdown* earlier in this document.

Probability: Low. This is only used when the system does not startup properly.

Impact: Medium. In the event the system does not start, there are other options to access the

system, including the IUCV terminal (discussed earlier in this section) and a maintenance virtual machine (discussed next). At least one of these methods should be implemented.

4. Maintenance virtual machine

Details: z/VM provides the ability to access and locally mount the root file system of one virtual machine within another. This enables one virtual machine to be used to repair another, in the event that files become corrupted or misconfigured which prevent normal startup.

Red Hat's Recommendation: Designate one guest as a maintenance virtual machine. This should be a relatively small guest that is powered off most of the time, but can be booted and used to fix problems in another. The maintenance virtual machine will primarily use the vmcp tool provided in Red Hat Enterprise Linux to link and attach one or more virtual DASD from one virtual machine to itself. This guest should contain basic tools to troubleshoot a system, such as LVM tools, a file system checker, and text editor.

Probability: Low. This is only required when a virtual machine does not startup normally.

Impact: Medium. Other options exist to access a non-booting system.

b) CPU and Memory

1. cpuplugd service

Details: The *cpuplugd* service is a daemon which continually monitors system load, then enables or disables virtual CPUs based on that load. When the system load drops below a threshold, additional virtual CPUs are disabled, thus improving z/VM application scheduling. When system load increases, virtual CPUs are enabled to handle the new workload.

Red Hat's Recommendation: Enable the *cpuplugd* service. This ensures that idle systems use less CPU resources, and busier systems get scheduled more often.

Probability: Medium. Depends on system load.

Impact: Medium. Recommended for workloads with peak utilizations throughout the day that are otherwise closer to idle.

2. Dynamic CPU add/remove

Details: Linux can enable and disable CPUs dynamically, also referred to as CPU hotplug. Using this, additional CPUs can be added to a virtual machine during high utilization, for example to complete workloads faster. They can then be disabled dynamically.

Red Hat's Recommendation: Configure virtual machines with the ability to perform CPU hotplug. See Appendix 3 for more information.

Probability: Low. Depends on workload. This is useful to manually manage CPU resources; however, z/VM may already do an adequate job of dynamically assigning and scheduling these resources.

Impact: Medium. Adding CPU resources can have a positive impact on virtual machine performance. Note, however, that this should not be used in conjunction with the *cpuplugd* service discussed above, as the effect is cancelled out.

3. /tmp file system

Details: The current SOE specifies a *tmpfs* file system for storing temporary files in /tmp, which

is a RAMdisk allocated out of real memory. Given the amount of files currently observed in /tmp, this can have a negative impact on system performance for Linux on System z, especially if VDISK is implemented.

Red Hat's Recommendation: Include /tmp in the existing LVM file system structure instead of using the *tmpfs* RAMdisk.

Probability: High. The current SOE for x86 specifies that /tmp use *tmpfs*.

Impact: High. Temporary files kept in real memory will waste shared resources in this virtual environment.

c) Networking

1. Increase MTU size

Details: Current MTU size is 1500 bytes. The z/VM VSWITCH can handle larger MTU sizes, which increase throughput and decrease CPU utilisation for applications that send packets larger than 1500 bytes.

Red Hat's Recommendation: Use MTU of 8992 bytes, which is the recommended size for z/VM VSWITCH.

Probability: Medium. Depends on application workload, and whether it benefits from sending large packets.

Impact: Medium. Can drastically increase throughput under certain network configurations.

2. Increase buffer count

Details: The *qeth* network driver supports larger internal buffers for virtual machines with heavy network load. One buffer consists of 16x4KB pages. The default value is 16 buffers, which occupies 1 MB of memory. The maximum value is 128, which occupies 8 MB of memory.

Red Hat's Recommendation: For virtual machines that are not memory constrained, and are expected to have heavy network throughput, increase buffers to 128. This is done by adding the following line to /etc/sysconfig/network-scripts/ifcfg-eth0:

```
OPTIONS=”buffer_count=128”
```

Probability: Low. Depends on workload, only designed for heavy network traffic.

Impact: High. Can significantly improve network bandwidth for a single virtual machine.

d) Shutdown

1. Change /etc/inittab behaviour

Details: z/VM provides the ability to trigger a *Ctrl+Alt+Del* with the **SIGNAL SHUTDOWN** command. By default, Linux handles *Ctrl+Alt+Del* with a reboot. However, z/VM assumes the signal should cause the guest to halt.

Red Hat's Recommendation: Change the behaviour of *Ctrl+Alt+Del* to shutdown instead of reboot. In /etc/inittab, make the following change:

```
ca::ctrlaltdel:/sbin/shutdown -t3 -h now
```

Probability: Low. Only triggered during SIGNAL SHUTDOWN, which occurs manually or during z/VM shutdown.

Impact: Low. z/VM could inadvertently reboot a system when the intent was to shutdown.

2. Configure shutdown actions

Details: By default when Linux is halted, the virtual machine remains logged on. This consumes unnecessary resources as z/VM must maintain the guest's memory allocation, along with any other resources in use.

Red Hat's Recommendation: Use the shutdown actions interface to set each virtual machine to logoff when powered off or halted. In /etc/rc.local, add the following lines:

```
/usr/sbin/chshut halt vmcmd logoff
/usr/sbin/chshut poff vmcmd logoff
```

Probability: High. The default action is to remain logged in after halt or power off.

Impact: Medium. Protects against a virtual machine remaining inadvertently logged on and consuming unused resources.

3.4. Standard Operating Environment

a) Installed Packages

Detail: There are a number of packages installed as part of the SOE that are either no longer needed due to the move to a virtualized System Z environment or no longer used.

Red Hat's Recommendation: Remove unnecessary packages and turn off unnecessary services, such as:

- Packages to be removed - OpenIPMI*, aspell, busybox, cyrus-sasl*, dmraid, ipsec-tools, iptables-ipv6, kudzu, nmap, system-config-network-tui, system-config-securitylevel-tui, kexec-tools, mdadm, dhcpv6-client, samba-common, subversion.
- Services to be turned off – netfs, rawdevices, xinetd (possibly required for Netbackup).

Probability: Low. Ensure the servers are kept up-to-date with relevant security patches.

Impact: Medium. By having unnecessary packages as well as services turned on, this could pose a security risk as well as an extra maintenance overhead.

Affected Environment: ALL

b) SELinux

Details: SELinux has been set to "permissive" on all of the servers. SELinux is the cornerstone of the Red Hat security model, and it should be enabled and enforcing at all times. Some organisations will not use SELinux for “internal” servers because it is believed that there is a minimal chance of compromise on those servers. Red Hat strongly recommends the use of SELinux for any Internet or customer facing servers.

Red Hat's Recommendation: Create appropriate SELinux profiles for any third party applications required on all servers and set SELinux to “enforcing” mode.

Probability: Medium. SELinux is part of Red Hat Enterprise Linux's security and we suggest it not be disabled.

Impact: Medium. SELinux enforces services security for Red Hat systems and services.

Affected Environment: ALL

c) iptables

Details: iptables are currently installed, but not configured on every server. Iptables provides an additional layer of security on a per system level, while still being able to be centrally managed by puppet.

Red Hat's Recommendation: Create appropriate iptables rules for any third party applications required on all servers and use puppet to centrally manage the rules on a per project or system purpose basis.

Probability: Medium. Iptables provides a host based firewall increasing each servers security .

Impact: Medium. With iptables on all Red Hat systems and services have an additional layer of security in front of them.

Affected Environment: ALL

d) SSH Configuration

Details: root logins via SSH are enabled in both the development and staging environments and X11 Forwarding is enabled in all environments. Disabling root logins in the other two environments ensures that login procedures and security are the same across all environments.

Whilst allowing X11 Forwarding is a potential weakness in security and should not be enabled in a blanket configuration for all systems in all environments and only be enabled where necessary [this sentence needs rewording]

Red Hat's Recommendation: root logins via ssh be disabled for all environments, not just production. X11 forwarding should be disabled by default and only enabled on systems where needed.

```
X11Forwarding no
PermitRootLogin no
```

Probability: Medium. This will not cause any issues, but will provide consistency and increase security across all environments.

Impact: Medium. Different procedures for the different environments can lead to mistakes and overall lowering in security .

Affected Environment: ALL

e) Centralised Logging

Details: For security reasons and administrative reasons, centralised logging should be considered. This will ensure logs to be available at a central location for debugging and security audit.

Red Hat's Recommendation: Using “rsyslog” instead of “syslog” remote logging feature should be enabled. Also, a centralised log server should be configured so all the server logs can be stored.

Probability: Low. This will not cause any issues, but this will help in reducing administrative overhead and increasing security.

Impact: Low. There will not be an immediate issue related to this. However, the centralised logging is a feature that should be considered.

Affected Environments: ALL

f) sysctl.conf / limits.conf

Details: Very limited optimisation and hardening has been done for both /etc/security/limits.conf and /etc/sysctl.conf. Doing further work in this area could prevent future resource problems, improve performance and increase security.

Red Hat's Recommendation: Further tuning and hardening of the limits.conf and sysctl.conf files on a per server purpose in sysctl.conf's case and per application in limits.conf's case should be performed.

Probability: Medium. This will not cause any issues, but could provide both performance and security improvements.

Impact: Low. There will not be an immediate issue related to this. However, tuning of these files is recommended for all systems.

Affected Environment: ALL

g) tmpfs Usage on /tmp and /var/tmp

Details: tmpfs is being used for both /tmp and /var/tmp which on virtual systems with tightly controlled and limited memory could potentially see RAM being used by files being stored in these directories.

Red Hat's Recommendation: /tmp and /var/tmp should not be mounted as tmpfs and instead be moved to use normal disk.

Probability: Medium. This will not cause any immediate issues, but a system could run out of memory if files are stored in these directories.

Impact: Medium. There will not be an immediate issue related to this. However, on limited memory virtual systems it's recommended this be done before problems occur in the future.

Affected Environment: ALL

h) auditd

Details: auditd is part of the Linux Auditing System and is configured with only the default configuration. A more comprehensive configuration would provide auditing capabilities such as who accessed or changed certain files on a system.

Red Hat's Recommendation: Further configuration of auditd can provide proper audit trail information for important system configuration files and/or application files, allowing the tracking and logging of who accessed or changed them.

Probability: Low. This will not cause any immediate issues, but this will help system auditing.

Impact: Low. There will not be an immediate issue related to this. However, further auditing will provide a way to find out who or what made changes to important files on a system.

Affected Environment: ALL

i) Kickstart

Details: The SOE kickstart file was created for x86 systems and contains certain packages that are no longer appropriate for a virtual system running on System Z or would be better moved to be controlled by puppet.

Red Hat's Recommendation: Many of the required packages should be moved out of the kickstart file and into puppet, with general system or multi-dependent packages being configured in packages-*.pp files and packages that are dependencies for certain applications being configured in that applications pp file as is currently being done, this will insure that puppet knows about all required packages, can restore them if necessary and that only required packages for a for the applications running on the system are installed.

In addition the @text-internet group is relatively large, only installing required packges is recommended.

Probability: Low. This will not cause any immediate issues, but will centralise configuration control further into puppet and lower the number of unnecessary packages on certain systems.

Impact: Low. There will not be an immediate issue related to this. However, what packages are installed and on which system is controlled from one place and with greater ease.

Affected Environment: ALL

j) Netbackup

Details: ~~Symantec Netbackup is currently configured to be installed on all systems. Netbackup is not the best backup solution for Red Hat Enterprise Linux on System z based systems. It has a history of causing performance issues that may result in problems with backups.~~

Red Hat's Recommendation: On System z Red Hat suggests to choose a backup solution that has a good track record and no known issues, e.g. IBM Tivoli Storage Manager (TSM).

Probability: Medium. Backups may not work properly and backups should be tested regularly for restore capability.

Impact: High. Missing or incomplete backups can result in the loss of data.

Affected Environment: ALL

3.5. Puppet Configuration

a) System arch based modules

Details: The move to using System Z has introduced an additional architecture that needs to be controlled and configured by puppet. Considerable work has been done in this area already, however further work needs to be done.

Red Hat's Recommendation: The creation of modules for all system architectures to allow for the splitting out of x86 specific software and configurations from the base-1_3 puppet module. For example the Netbackup puppet configuration, which is not the ideal backup solution on System Z based systems, but is currently located in base-1_3.

Probability: Medium. Inappropriate packages and configuration for System Z based virtual machines will be installed because they are located in base and not a x86 based arch module.

Impact: Medium. If inappropriate packages are installed, such as Netbackup, expected results and functions might not happen.

Affected Environments: ALL

b) Project specific configuration in base module

Details: There is still project specific configuration located in the base-1_3 puppet module.

Red Hat's Recommendation: The removal of any project specific configuration still left in base-1_3, such as the configuration in the network-common.pp. Moving this to the Acme module would consolidate all the project configuration into one place as well as simplify the puppet configuration in base-1_3.

Probability: Low. This will not cause any immediate issues, but having standardised locations for all project configuration makes management easier.

Impact: Low. There will not be an immediate issue related to this. But will provide a simplified puppet configuration.

Affected Environments: ALL

c) General puppet configuration

Details: Large and complicated puppet configurations can place considerable system load on the puppet master server. If puppets configuration is broken down into as many small and simple modules as possible and then configured only on the systems that require them, the puppet master has a lot less work to do and a higher performance can be extracted from the system before load becomes an issue.

Red Hat's Recommendation: The base module should be kept as small and as simple as possible, to keep system load on the puppet master within manageable levels as the puppet configurations and use grow over time.

Probability: Low. This will not cause any immediate issues, but something to plan for in the future.

Impact: Low. There will not be an immediate issue related to this. The puppet master can place considerable load on a system in large and complicated environments.

Affected Environments: ALL

d) Minor issues

Details: The additional templates and configuration files for limits.conf and stsctl.conf appear to be missing.

Postfix and puppet pp files are included twice in the base-1_3 init.pp file.

Red Hat's Recommendation: Create the required template and configuration files for the management of limits and sysctl so they will work as expected.

Remove the duplicate entries from the init.pp in base-1_3 puppet module.

Probability: Medium. This will not cause any immediate issues, but the expected configurations for sysctl and limits will not be in place and unexpected results at both the OS and application level might occur.

Impact: Medium. There will not be an immediate issue related to this. However once load testing is performed, applications might reach the defaults system limits and behave in an unexpected manner.

Affected Environments: ALL

3.6. Appendix

a) Starting up scripts

The following example shows REXX scripting to present a menu at each virtual machine login. This should be added to the user linuxadm's PROFILE EXEC. This menu is only shown during an interactive login on the 3270 console. If the user is started automatically via the XAUTOLOG command, the menu is bypassed and the guest is booted with IPL. Note that after adding these lines, it is recommended to start the virtual machine automatically when z/VM starts by adding XAUTOLOG statements in the user autolog1's PROFILE EXEC.

```
00006 'PIPE CP QUERY' userid() '| var user'
00007 parse value user with id . dsc .
00008 iplDisk = 200
00009 if (userid() <> 'LINUXADM') then
00010 do
00011 vmfclear
00012 say ' '
00013 say ' '
00014 say 'Press <Enter> to boot normally'
00015 say ' '
00016 say ' '
00017 say 'MENU'
00018 say '----'
00019 say '1. Boot Linux (Default)'
00021 say '2. Start interactive session with CMS'
00022 say '3. Logout'
00023 say ' '
00024 say ' '
00025 say 'Enter Choice'
00026 'CP TERM MORE 0 0'
00027 if (dsc = 'DSC') then /* user is disconnected */
00028   'CP IPL' iplDisk
00029 else  /* user is interactive -> prompt */
00030 do
00031   parse upper pull answer .
00032   if (answer = '1') then
00033     'CP IPL' iplDisk
00034   else if (answer = '2') then
00035   say ' '
00036   else if (answer = '3') then
00037     'CP LOG HOLD'
00038   else
00039     'CP IPL' iplDisk
00040 end /* else */
00041 end
```

b) TTY Terminal over IUCV

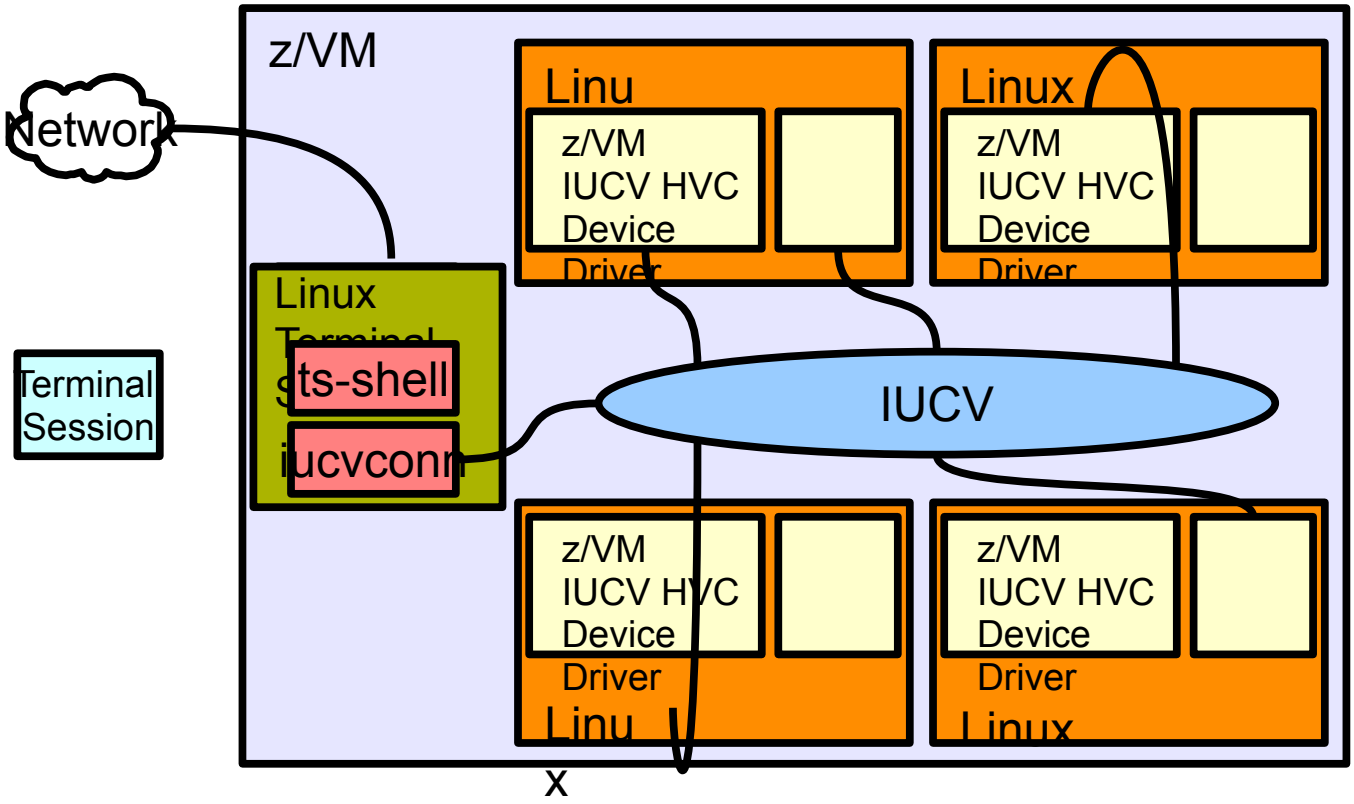

c) Dynamic CPU

Dynamic CPU configuration

(from https://access.redhat.com/kb/docs/DOC-2399)

CPU resources are controlled in two places. First in the profile of the guest, and second in the zipl bootloader. Follow these steps to add CPU resources:

1. Using a 3270 emulator, connect to z/VM using the MAINT user-id
2. Using xedit, edit the user-id DIRECT as follows: xedit USER DIRECT
3. Find the stanza which describes the default user profile for the Red Hat Enterprise Linux guest
4. Edit the line MACHINE ESA X to reflect the correct number of CPUs
5. Add lines as necessary below to reflect the CPUs. For example:

```
PROFILE RHDFLT
MACHINE ESA 4
CPU 00 BASE
CPU 01
CPU 02
CPU 03
```

6. Save by issuing the FILE command, then logoff
7. Log onto the Red Hat Enterprise Linux guest as root and edit /etc/zipl.conf
8. Change the entry for the kernel, specifically, the line which reads parameters. It should read parameters="root=/dev/VolGroup00/LogVol00". Change the line to read parameters="root=/dev/VolGroup00/LogVol00 maxcpus=4". **Note:** setting maxcpus to a value LESS than the value of maxcpus still adds the CPUs to the system, but leaves them unused by the Red Hat Enterprise Linux guest.
9. Run the command zipl to apply the changes and reboot. The CPU resources will now be available to the Red Hat Enterprise Linux guest.

Dynamic CPU Enablement

(from https://access.redhat.com/kb/docs/DOC-2406)

Enabling CPU resources is straightforward. Simply follow these steps:

1. Run the command cd /sys/devices/system/cpu
2. This directory lists the CPUs available to the running guest.
3. To query the status of the CPUs, issue the command cat cpuX/online where X refers to the cpu-id. If the value returned is 1, then that CPU is online and in use. If the value returned is 0 then that CPU is unused and available.
4. Unused CPUs can be brought online by issuing the command echo "1" > cpuX/online
5. Conversely, CPUs can be taken offline by issuing the command echo "0" > cpuX/online