

Virtualization Standards for Business Continuity: Part 8

This is the eighth of a series of articles defining the policies, guidelines, standards, and procedures that provide the foundation of a virtualized environment, thus enabling business continuity, disaster recovery, and high availability, with an emphasis toward Return On Investment (ROI).

The focus of this article will be to describe a variety of VIO server settings and requirements to support the business continuity methodology described by this series of articles. In a previous article, a script named “mkviolpar” was provided for creating VIO Server LPAR's, this article will describe the procedures for installing a VIO Server operating system onto an LPAR created by the “mkviolpar” script, and a number of configuration parameters associated with the VIO Server.

The “mkviolpar” uses default values for many of the VIO server configuration parameters, however these can be modified by the user on an as needed basis. These parameters include settings for CPU, memory, physical and virtual resource sharing. Before modifying the values for the CPU settings, the users should know and understand some basic Power5 and Power6 architecture. The CPU settings are divided into two categories, processing units and virtual processors.

Processing units represent the amount of physical CPU processing power that can be utilized by the VIO server LPAR. The minimum amount of CPU processing units that can be assigned to an LPAR is 0.10 or 1/10th of a processor, and can be incremented above that in increments of 0.01 processing units or 1/100th of a processor. A single virtual processor is represented on an LPAR as two active logical processors on a system with symmetric multi-threading (SMT) enabled, or a single active logical processor if SMT is not enabled. Each logical processor translates into a processing thread. The reason this is important to know is the minimum number of processing units that can be assigned to a virtual processor is 0.10 or 1/10th of a processor. This means that for every virtual processor required by an LPAR, 0.10 processing units must also be allocated. So an LPAR requiring four (4) virtual processors (8 logical processors or SMT threads) requires a minimum of 0.40 processing units.

IBM typically recommends 1.00 processing units be assigned for every gigabit of network bandwidth provided by the VIO server LPAR. Experience has shown this is not normally necessary, usually a minimum value for the number of processing units can be set to 0.20 with a minimum of 2 virtual processors. The maximum network throughput of a VIO server has a hard limit of about 3.6 gigabits, so the absolute maximum amount of processing units that should ever be assigned to a VIO server is 3.6. The maximum number of virtual processors should be determined by the number of I/O adapters and the load being serviced by the VIO server. Again, experience has shown the maximum CPU values on a VIO server can normally be set to 2.00 processing units, with four (4) virtual processors.

The desired values for processing units and virtual processors on a VIO server can be determined by the system administrator based on the expected processing load on the entire frame. If processing power is expected to be in short supply, it may be best to preallocate CPU on the VIO servers to higher values to ensure availability. If CPU requirements are expected to fluctuate on the frame, the desired CPU values may be set to low values, and then allow the hypervisor to control the allocation and deallocation of processing power. Table 7.1 provides a list of recommended initial CPU parameters for a VIO Server.

<i>Description</i>	<i>mkviolpar Variable</i>	<i>Recommended Value</i>
Minimum Virtual Processors	MIN_PROCS	2
Desired Virtual Processors	DESIRED_PROCS	2

Virtualization Standards for Business Continuity: Part 8

<i>Description</i>	<i>mkviolpar Variable</i>	<i>Recommended Value</i>
Maximum Virtual Processors	MAX_PROCS	4
Minimum Processing Units	MIN_PROC_UNITS	0.20
Desired Processing Units	DESIRED_PROC_UNITS	0.20
Maximum Processing Units	MAX_PROC_UNITS	2.00
Processing Units Sharing Mode	SHARING_MODE	ON
Processing Units Uncapped Weighting	UNCAP_WEIGHT	128

Table 7.1: CPU Parameters for VIO Server

The memory requirements of a VIO server is usually an easy assignment. The values shown in Table 7.2 for the VIO server memory parameters work very well in most scenarios.

<i>Description</i>	<i>mkviolpar Variable</i>	<i>Recommended Value</i>
Minimum Memory	MIN_MEM	1024 Mb
Desired Memory	DESIRED_MEM	1024 Mb
Maximum Memory	MAX_MEM	2048 Mb

Table 7.2: Memory Parameters for VIO Server

If a large number of networks or SAN storage will be provided by the VIO servers, the memory values may need to be increased.

One of the underlying themes of this business continuity methodology is that all installation and update processes should be performed through a centralized location, such as a Network Installation Manager (NIM) Server. The installation of the VIO Server can be performed via removable media such as CD-ROM or DVD, however the preferred method is to upload the media to the NIM server and perform all installations from that location. Enabling this installation procedure requires the VIO server media be uploaded in it's entirety to a location on the NIM server. The following procedure assumes the CD-ROM or DVD media contains VIO Server version 1.3.0.1, if your media is a different version, modify the procedure as necessary to reflect your version number. Copy the VIO server CD or DVD to the NIM server where the location the media will be copied to is `"/export/lpp_source/vio/vio_1301"`.

```
mkdir -p /cdrom
mkdir -p /export/lpp_source/vio/vio_1301
mount -v cdrfs -o ro /dev/cd0 /cdrom
cd /cdrom
find . -print | cpio -puvd /export/lpp_source/vio/vio_1301
```

Once the VIO media is copied to the NIM Server, the VIO server can be created using the following

Virtualization Standards for Business Continuity: Part 8

automated installation method from the NIM server. The following example uses VIO Server host names and TCP/IP address values from previous articles in this series:

```
D="-d /export/lpp_source/vio/vio_1301" # VIO Server media location on NIM Server
H="-h 10.1.2.123" # IP address of HMC for password-less ssh access
S="-s Server-9119-590-SN12A345B" # Managed System name on which to build VIO Server
P="-p dalapvio00" # VIO Server LPAR Name (assumed to already exist)
R="-r dalapvio00" # VIO Server Profile Name (assumed to already exist)
I="-i 10.1.2.50" # IP Address of VIO Server to build
M="-S 255.255.255.0" # Network Mask of VIO Server to build
G="-g 10.1.2.254" # Network Default Gateway of VIO Server to build
B="-P auto" # Network Speed Setting
C="-D auto" # Network Duplex Setting
N="-n" # Configure Client network after install
A="-N" # Accept I/O Server License Agreement
/usr/sbin/installios -f -u # Cleanup any previous failed installation attempts
/usr/sbin/installios ${D} ${H} ${S} ${P} ${R} ${I} ${M} ${G} ${B} ${C} ${N} ${A}
```

<i>VIO Server Name</i>	<i>VIO Server Location</i>	<i>Frame Name</i>
<i>dalapvio00</i>	First VIO Server node on the 1st frame	Server-9119-590-SN12A345B
<i>dalapvio01</i>	Second VIO Server node on the 1st frame	Server-9119-590-SN12A345B
<i>dalapvio02</i>	First VIO Server node on the 2nd frame	Server-9119-590-SN67D890E
<i>dalapvio03</i>	Second VIO Server node on the 2nd frame	Server-9119-590-SN67D890E

Table 7.2: Example VIO Server Host Names and Frame Names from Previous Article

Once the VIOS operating system has been installed on a VIO LPAR, the I/O adapters must be configured to be shared out to the client LPARs. The exact configuration of each VIO server will depend upon the quantity and characteristics of the physical adapters and the storage devices used. The following example configuration assumes each VIO server has four (4) ethernet adapters to be configured for redundant access to two(2) different networks, and 2 fibre channel adapters per VIOS providing MPIO access to storage. Your configuration will likely be different, this example is provided only as a guide to scripting your configuration. *It is assumed these example commands would be executed from the "oem_setup_env" prompt on each VIO server:*

Virtualization Standards for Business Continuity: Part 8

```
for i in 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
do
    ifconfig en${i} down
    ifconfig en${i} detach
    ifconfig et${i} down
    ifconfig et${i} detach
    rmdev -Rdl ent${i}
    rmdev -Rdl en${i}
    rmdev -Rdl et${i}
done
# Rediscover all ethernet adapters
cfgmgr
# Turn off the chksum_offload parameter on all adapters
for i in 10 9 8 7 6 5 4 3 2 1 0
do
    lsattr -El ent${i} | grep chksum_offload
    chdev -l ent${i} -a chksum_offload=no
    chdev -l ent${i} -a chksum_offload=no -P
    lsattr -El ent${i} | grep chksum_offload
done
# create an etherchannel using physical adapters ENT0 and ENT2
unset S T A B N
S="-s pseudo"
T="-t ibm_ech"
A="-a adapter_names=ent0,ent2"
N="-a netaddr=10.1.2.0"
if mkdev -c adapter ${S} ${T} ${A} ${B} ${N}
then
    /usr/lib/methods/defif
fi
# create an etherchannel using physical adapters ENT1 and ENT3
unset S T A B N
S="-s pseudo"
T="-t ibm_ech"
A="-a adapter_names=ent1,ent3"
N="-a netaddr=10.1.3.0"
```

Virtualization Standards for Business Continuity: Part 8

```
if mkdev -c adapter ${S} ${T} ${A} ${B} ${N}
```

```
then
```

```
  /usr/lib/methods/defif
```

```
fi
```

```
lsdev -Cc adapter | grep ent
```

```
chdev -l ent8 -a mode=round_robin
```

```
chdev -l ent9 -a mode=round_robin
```

```
lsattr -El ent8
```

```
lsattr -El ent9
```

```
# build the SEA's (shared ethernet adapters)
```

```
S="-sea ent8"
```

```
V="-vadapter ent4"
```

```
D="-default ent4"
```

```
I="-defaultid 2"
```

```
/usr/ios/cli/ioscli mkvdev ${S} ${V} ${D} ${I}
```

```
S="-sea ent9"
```

```
V="-vadapter ent6"
```

```
D="-default ent6"
```

```
I="-defaultid 3"
```

```
/usr/ios/cli/ioscli mkvdev ${S} ${V} ${D} ${I}
```

```
/usr/ios/cli/ioscli lsmmap -all -net
```

If TCP/IP addresses are to be assigned to the VIO servers and the VIO servers were built using the “mkviolpar” script, then the addresses should be configured on the unused virtual ethernet adapters. Notice in the above example virtual ethernet adapters “ent4” and “ent6” were used to configure the SEA adapters, this leaves “ent5” and “ent7” as unused. Also observe in the VIO server configuration these adapters point to the VLAN ID's used by “ent4” and “ent6” respectively. So any network traffic traveling across these adapters (ent5 or ent7) will be redirected to the VLAN ID's associated with network adapters “ent4” and “ent6”. The VIO server TCP/IP addresses should be placed on the “ent5” and/or “ent7” virtual adapters so they can be modified and manipulated as necessary without interrupting the client LPAR network stack of adapters.

The next section of code provides a scripting mechanism to configure the VIO server storage adapters and disks to enable them to be shared out to client LPAR's. This example configuration assumes each VIO server has two(2) fibre channel adapters providing multi-path I/O (MPIO) to each LUN and that a quantity of LUN's have been assigned to the fibre channel adapters. This should be used only as a

Virtualization Standards for Business Continuity: Part 8

guide for creating a script specific to your VIO server configuration. *It is assumed these example commands would be executed from the "oem_setup_env" prompt on each VIO server:*

```
for i in 0 1
do
  chdev -l fcs${i} -a init_link=pt2pt -P
  chdev -l fscsi${i} -a dyntrk=yes -P
  chdev -l fscsi${i} -a fc_err_recov=fast_fail -P
done

cfgmgr -vl fcs0
cfgmgr -vl fcs1

# Change reserve policy on hdisks to "no_reserve"
for i in $(lsdev -Cc disk -s fcp -F name )
do
  chdev -l ${i} -a reserve_policy=no_reserve -P
  chdev -l ${i} -a algorithm=round_robin -P
  chdev -l ${i} -a hcheck_interval=20 -P
  chdev -l ${i} -a queue_depth=16 -P
  chdev -l ${i} -a pv=yes -P
done
```

The Business Continuity policies, guidelines, standards, and procedures to be learned from this article are as follows:

Policies: *Those tasks that must be implemented*

- Single points of failure will be eliminated in the VIO server design.
- Dual VIO servers with standardized configurations will be implemented for all implementations.

Guidelines: *Those tasks that should be implemented*

- VIO Servers should be allocated an adequate amount of CPU and memory to provide the performance necessary to support the expected client LPAR load. The default values should be calculated for the network and storage communication load across all environments.
- Network access through the VIO servers should be configured as best suits the business functions utilizing the bandwidth. This may include the use of etherchannels or may not. Whatever mechanism is selected, it should be used as the standard build across all VIO servers

Virtualization Standards for Business Continuity: Part 8

for consistency, repeatability, and ease of training new personnel.

Standards: *Technical specifications derived from the policies and guidelines*

- This article provides an initial CPU and memory estimate for configuring VIO Servers.
- This article provides a standard physical configuration for network adapters.
- This article provides a standard physical configuration for fibre channel adapters.

Procedures: *Step-by-step implementation instructions of the standards*

- This article provides a procedure for changing the default CPU and memory values utilized by the “mkviolpar” script.
- This article provides a procedure for configuring physical network adapters for use by the VIO servers.
- This article provides a procedure for configuring physical fibre channel adapters and storage LUN's for use by the VIO servers.

The next article in this series will discuss standardized procedures for creating client LPAR's and configuring the resources associated with each. These resources will include CPU, memory, naming standards, slot numbering, network, and storage adapters.

Dana French
President, Mt Xia, Inc.
<http://www.mtxia.com>
615.556.0456