

# VIOS Performance Monitoring Setup Instructions

December 13, 2016

Please send questions to: [idoctor@us.ibm.com](mailto:idoctor@us.ibm.com)

## Summary

These steps allow you to automatically collect NPIV, NMON and SEA stats from VIOS partitions and send the data to IFS directories located on an IBM i partition. There is a monitor that runs on the IBM i partition which imports the IFS files into database file members that can be analyzed with the iDoctor client. Retention periods are set for both the IFS and database file members. Requirements and setup instructions are included below.

## Requirements

1. Install the latest iDoctor client and server builds located on our website. Note that the iDoctor server build will need to be updated on the partition that the VIOS data is sent to and analyzed on. The iDoctor client will need to be updated on any PC that will be used to work with the data.

<http://public.dhe.ibm.com/services/us/igsc/idoctor/html/downloadOptions.html>

2. Unpack the .tar file `idoctor.tar` on the VIOS partitions that you will be collecting data from. See Section 1 – Installing iDoctor Scripts.
3. Set up ssh between the VIOS and IBM i. See Section 2 - SSH Setup Instructions.
4. Set up the VIOS Monitor on the VIOS partition(s) and the target data directories on the IBM i Analysis System. See Section 3 - VIOS Monitor Setup Instructions.

5. Set up the Directory Monitor on the IBM i Analysis System. See Section 4 – Directory Monitor Setup instructions.

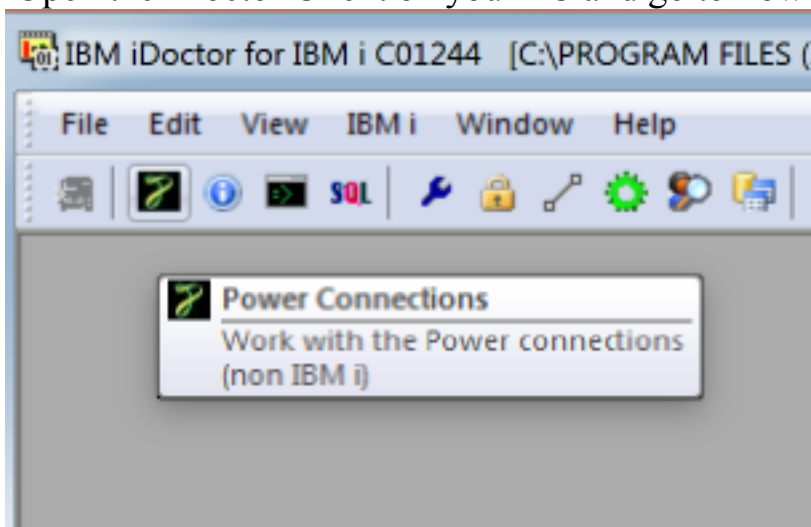
## Section 1. Installing iDoctor Scripts

There is a file a file that exists as part of the iDoctor client build located on your PC that contains the iDoctor scripts to install on the VIOS.

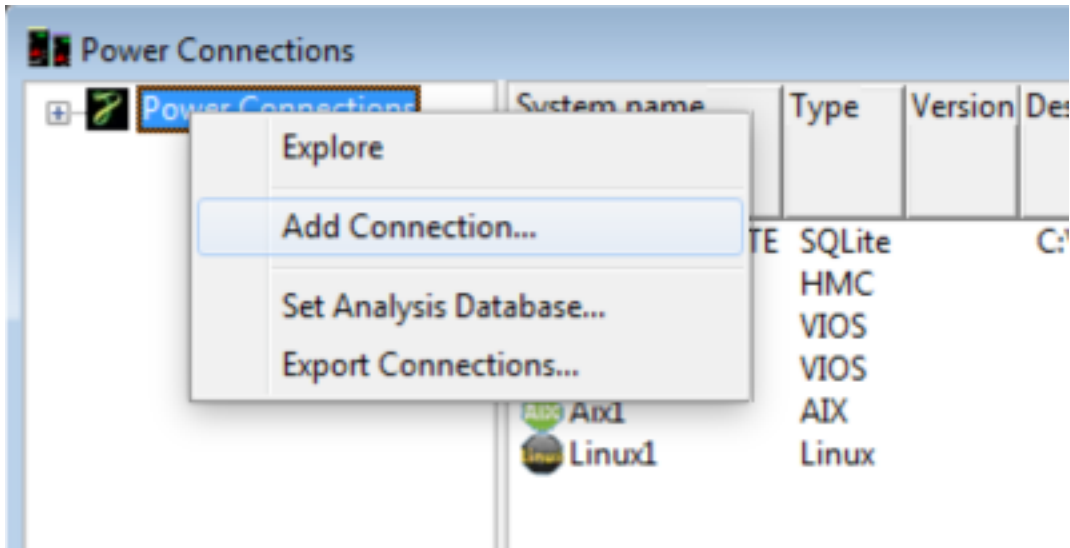
C:\Program Files (x86)\IBM\iDoctor\ServerUpdates\AIX\idoctor.tar

This can be done over an SSH connection from the PC to the VIOS or when connecting to the VIOS with iDoctor GUI for the first time the files will be transferred and unpacked if they are not already there. Here we show how to connect to the VIOS using the iDoctor GUI. If sending the file manually using SSH the file should be unpacked into /tmp/idoctor.

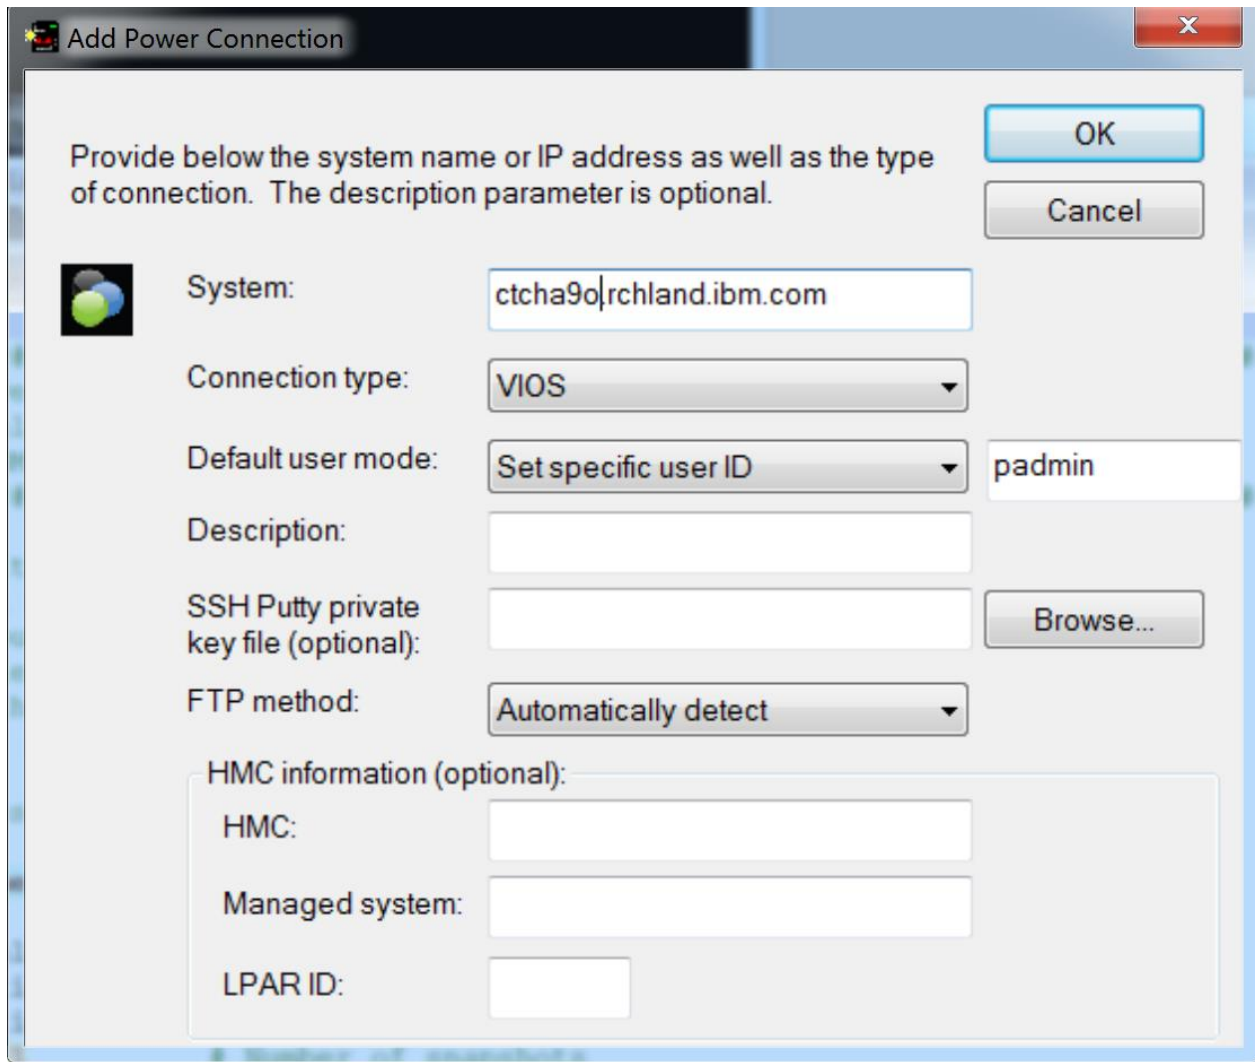
1. Open the iDoctor Client on your PC and go to Power Connections.



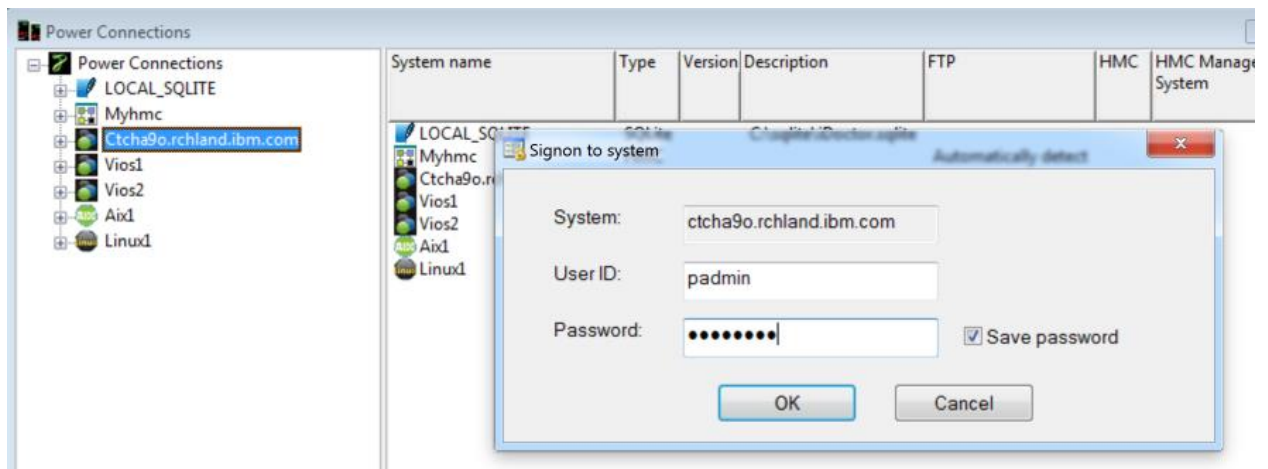
2. Right click on Power Connections and select 'Add Connection'.




3. Fill out the 'Add Power Connections' window as below with the fully qualified host name or IP address of your VIOS partition. Note that padmin must be used for the user id. Click ok.




4. Click on the VIOS partition in the left pane, enter your login credentials, and click ok.



- Click yes when the following window appears asking if you would like to install the scripts.



- The following messages in the Remote Command dialogue indicate successful installation of the scripts.



## Section 2. SSH Setup Instructions

The SSH Setup Instructions below will walk you through setting up SSH between the VIOS and IBM i.

- On the VIOS, generate the public/private key pair for ssh.
  - After logging in as padmin enter oem\_setup\_env
  - cd /home/padmin
  - ssh-keygen -t rsa -f id\_rsa -N "

```
# cd /home/padmin
# ssh-keygen -t rsa -f id_rsa -N ''
Generating public/private rsa key pair.
Your identification has been saved in id_rsa.
Your public key has been saved in id_rsa.pub.
The key fingerprint is:
ea:65:af:41:86:6d:8c:f1:c7:cf:ed:bb:97:9d:43:4d root@ctcvha9e
The key's randomart image is:
+--[ RSA 2048]-----+
|
|
|   .           |
|  B .         E|
|  o S o       ..|
|   = . o . . .|
|  . +   o o +  |
|  . o o   . +o |
|  . . .   ++.  |
+-----+

```

2. Determine which user will be used for ssh connections to the IBM i partition. Log on to the IBM i partition and ensure the user has a /home/myuserid directory created where myuserid is the user that will be used for ssh connections.
  - a. qsh
  - b. mkdir /home/myuserid/.ssh (if the directory does not exist)
  - c. touch /home/myuserid/.ssh/authorized\_keys (if the authorized key file does not exist)
  - d. ls -la to view permissions. Permissions may need to be changed in a later step.

```
QSH Command Entry
$
> mkdir /home/BSMENGENES/.ssh
$
> ls -la
total: 28 kilobytes
drwxrwsrwx  3 BSMENGENES  0
drwxrwsrwx  7 QSYS        0
drwxrwsrwx  2 BSMENGENES  0
```

3. Ensure the user profile has the home directory specified. Use DSPUSRPRF to check and CHGUSRPRF to change the profile if necessary.

```

Change User Profile (CHGUSRPRF)

Type choices, press Enter.

Locale . . . . . *SAME
-----
User options . . . . . *NONE          *SAME, *NONE, *CLKWD...
      + for more values
-----
User ID number . . . . . 131          1-4294967294, *SAME
Group ID number . . . . . *NONE          1-4294967294, *SAME, *GEN...
Home directory . . . . . /home/BSMENGES
-----

```

4. FTP to the VIOS partition to retrieve the public key file or have it sent to an intermediate system if FTP is not allowed. Note that ASCII transfer should be used or the key file should be zipped prior to sending. If FTPing from the IBM i to VIOS:
  - a. ftp to the vios
  - b. ascii
  - c. cd /home/padmin
  - d. namefmt 1
  - e. get id\_rsa.pub /home/myuserid/.ssh/id\_rsa.pub
  - f. quit

```

> ascii
200 Type set to A; form set to N.
> cd /home/padmin
250 CWD command successful.
> namefmt 1
500 'SITE NAMEFMT 1': command not understood.
Client NAMEFMT is 1.
> get id_rsa.pub /home/bsmenges/.ssh/id_rsa.pub
229 Entering Extended Passive Mode (|||56151|)
150 Opening data connection for id_rsa.pub (395 bytes).
226 Transfer complete.
396 bytes transferred in 0.173 seconds. Transfer rate 2.291 KB/sec.

```

5. Ensure the IBM i is configured properly for SSH to work. Below is a good link which summarizes what is needed.

[http://www-304.ibm.com/partnerworld/wps/servlet/ContentHandler/pw\\_com\\_porting\\_tools\\_openssh](http://www-304.ibm.com/partnerworld/wps/servlet/ContentHandler/pw_com_porting_tools_openssh)

Key points:

- Ensure 5733-SC1 -- IBM portable utilities for i is installed on the system.
- The user's home directory must not have public write authority ( `chmod go-w /home/myuserid` )
- The following directories for the user must not have any public authorities.  
`/home/myuserid/.ssh`
- `/home/myuserid/.ssh/authorized_keys` authorities
- ( `chmod go-rwx /home/myuserid/.ssh` and `chmod go-rwx /home/myuserid/.ssh/authorized_keys` )

```
> chmod go-w /home/bsmenges
$
> chmod go-rwx /home/bsmenges/.ssh
$
> chmod go-rwx /home/bsmenges/.ssh/authorized_keys
$
```

6. The public key must be located in the `/home/myuserid/authorized_keys` file. Use the `cp` command to copy the `id_rsa.pub` file into the `authorized_keys` file. Only use the `cp` command if the file is empty or you do not care about its contents as `cp` will overwrite the file. If the `authorized_keys` file already contains data, keys from other systems for example, use the `cat` command instead as follows. If you will be having multiple VIOS send data to the same IBM i, then you will need to bring down the key files one at a time and use `cat` to add them to the `authorized_keys` file.

**Example using `cp` (data in the file will be replaced) `cp id_rsa.pub authorized_keys`**

```
> cp id_rsa.pub authorized_keys
$
> ls
authorized_keys id_dsa.pub id_rsa.pub
```

**Example using `cat >>` (data will be appended to the end of the file)**

**`cat id_rsa.pub >> authorized_keys`**

```
==> cat id_rsa.pub >> authorized_keys
```



7. Ensure that the ssh daemon is in listen mode. Netstat option 3 and look for local port ssh. Use STRTCPSVR SERVER(\*SSHD) to start the server. It is also recommended to set this server to autostart when TCP is started on the system. This can be done in iNav or by updating the QUSRSYS/QATOCSTART.

```
Work with IPv4 Connection Status                                     System:  ID0C710
Type options, press Enter.
 3=Enable debug  4=End  5=Display details  6=Disable debug
 8=Display jobs

  Opt  Remote      Remote      Local      Idle Time  State
  ---  -
  *    *          *          ftp-con > 113:16:04 Listen
  *    *          *          ssh        000:00:35 Listen
```

8. Now that the IBM i setup is complete, SSH from the VIOS partition to the IBM i in order to add the IBM i to the Known Hosts File on the VIOS.

- a. After logging in to the VIOS issue the command `ssh myuserid@myipaddress` where myuserid is the userid configured on the IBM i partition and my IP address is the IP address or hostname.domain name of the IBM i partition. If using hostname VIOS must be able to perform DNS lookup.
- b. Answer yes when prompted. The IBM i will be added to the list of known hosts.
- c. Log on with the password for myuser.
- d. Log off the IBM i partition by typing exit and press enter.

```
login as: padmin
padmin@ctcvha9e.rchland.ibm.com's password:
Last unsuccessful login: Wed Jan 21 15:16:30 CST 2015 on ssh from 9.10.75.129
Last login: Sat Mar 21 09:04:29 CDT 2015 on ftp from idoc710.rchland.ibm.com
$ oem_setup_env
# ssh bsmenges@9.5.68.31
The authenticity of host '9.5.68.31 (9.5.68.31)' can't be established.
ECDSA key fingerprint is cb:33:f9:b8:3f:fd:64:34:69:fa:e0:1e:9f:05:d9:1f.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '9.5.68.31' (ECDSA) to the list of known hosts.
bsmenges@9.5.68.31's password:
$ █
```

10. Now that everything is set up on the VIOS and IBM i partitions, test that you are able to logon without a password.

- a. `ssh -i /home/padmin/id_rsa myuserid@myIPaddress`
- b. If you get a \$ prompt without further prompting or messages, you are logged on the IBM i. Type exit, and enter to log off.
- c. If you are prompted for a password, then something is wrong in the configuration or with the ssh daemon on the server side. Try ssh again with -iv, -ivv, or -ivvv flags for various levels of logging that will help to pinpoint the error (verbose, very verbose, or very very verbose). The sshd logs on the IBMi side can also be checked for errors.

### Section 3. VIOS Monitor Setup Instructions

This section will walk you through setting up data directories on the IBM i partition and cron scheduler entries on the VIOS that will automatically collect and send NMON, NPIV, and SEA data to the IBM i.

1. Create the data directories on the IBM i. In this example we create three sub directories in the user's home directory. One for each for NMON, NPIV, and SEA data. If you will be sending from multiple VIOS you can send all of the data to the same directories, create one directory for each VIOS, or set it up however you want. In the next section, the setup for the Directory Monitor will allow you to configure the monitor with the

various data locations for import by VIOS host name, data type, and target library.

- a. qsh
- b. mkdir /home/myuserid/nmon
- c. mkdir /home/myuserid/npiv
- d. mkdir /home/myuserid/sea

```
$
> mkdir /home/bsmenges/sea
$
> mkdir /home/bsmenges/npiv
$
> mkdir /home/bsmenges/nmon
$
```

2. Enable SEA Advanced Accounting. The 'seastat' command is used to collect samples from SEA adapters on the VIOS. Accounting needs to be enabled for each adapter in order for the command to collect data.
  - a. Enter 'oem\_setup\_env' followed by 'lsdev -Cc adapter' to get a list of all adapters on the VIOS. To filter out just the SEA adapters use 'lsdev -Cc adapter | grep Shared'. Note that 'grep' is a case sensitive search.
  - b. For each SEA, enter 'chdev -dev entxx -attr accounting=enabled' replacing entxx with the name of the SEA adapter. Repeat this for all SEA adapters you want to monitor. Running the command with 'accounting=disabled' turns off advanced accounting.
3. Add cron schedule entries on the VIOS. The following steps are using crontab -e to modify the cron file manually using the vi editor. This can be tricky if it is your first time using vi. There is a good vi cheat sheet located here: <http://www.lagmonster.org/docs/vi.html>. Each time you modify the file you will use the following process.
  - a. Log on to the VIOS with padmin.
  - b. Type oem\_setup\_env and press enter.
  - c. Type crontab -e and press enter.
  - d. Make changes to the file.
  - e. :x and press enter to save the changes.

4. Below is a screenshot showing the required entries in cron. An explanation of the entries follow.

```
-e
0 * * * * /tmp/idoctor/nmon_monitor.sh 300 12
* * * * * /tmp/idoctor/scp_monitor.sh /home/padmin/id_rsa bsmenges 9.5.65.1 /home/bsmenges/nmon nmon
0 * * * * /tmp/idoctor/npiv_monitor.sh 13 300
* * * * * /tmp/idoctor/scp_monitor.sh /home/padmin/id_rsa bsmenges 9.5.65.1 /home/bsmenges/npiv npiv
0 * * * * /tmp/idoctor/sea_monitor.sh 13 300
* * * * * /tmp/idoctor/scp_monitor.sh /home/padmin/id_rsa bsmenges 9.5.65.1 /home/bsmenges/sea sea
```

**a. NMON Monitor:**

```
0 * * * * /tmp/idoctor/nmon_monitor.sh 300 12
```

This starts a new NMON collection with 5 minute intervals every hour (300 seconds x 12 intervals).

**b. SCP Monitor for NMON data:**

(\* \* \* \* \* indicates the entry will run every minute and check if anything is ready to be sent)

```
/tmp/idoctor/scp_monitor.sh /home/padmin/id_rsa myuserid  
myIPaddress /home/myuserid/nmon nmon
```

myuserid: User ID configured for SSH on the IBM i system.

MyIPaddress: IP address of the IBM i system.

/home/myuserid/nmon: Remote directory where nmon data for this system should be sent.

Nmon: Tells the SCP Monitor that this is nmon data.

**c. NPIV Monitor:**

```
0 * * * * /tmp/idoctor/npiv_monitor 13 300
```

This starts a new NPIV monitor every hour. 13, 5 minute snapshots are collected.

This results in a full 60 minutes (12, 5 minute intervals) of data each hour.

**d. SCP Monitor for NPIV data:**

(\* \* \* \* \* indicates the entry will run every minute and check if anything is ready to be sent)

```
/tmp/idoctor/scp_monitor.sh /home/padmin/id_rsa myuserid  
myIPaddress /home/myuserid/npiv npiv
```

myuserid: User ID configured for SSH on the IBM i system.

MyIPAddress: IP address of the IBM i system.

/home/myuserid/npiv: Remote directory where npiv data for this system should be sent.

npiv: Tells the SCP Monitor that this is nmon data.

**e. SEA Monitor**

```
0 * * * * /tmp/idoctor/sea_monitor 13 300
```

This starts a new SEA monitor every hour. 13, 5 minute snapshots are collected.

This results in a full 60 minutes (12, 5 minute intervals) of data each hour.

**f. SCP Monitor for SEA data:**

(\* \* \* \* \* indicates the entry will run every minute and check if anything is ready to be sent)

```
/tmp/idoctor/scp_monitor.sh /home/padmin/id_rsa myuserid  
myIPAddress /home/myuserid/sea sea
```

myuserid: User ID configured for SSH on the IBM i system.

MyIPAddress: IP address of the IBM i system.

/home/myuserid/sea: Remote directory where sea data for this system should be sent.

sea: Tells the SCP Monitor that this is sea data.

5. Example of how to enter the above entries using crontab –e.

a. crontab –e

b. Press the ‘o’ key to open a new line after the cursor. This places the editor into insertion mode.

c. Type in the following (NMON monitor from above is used as an example): 0 \* \* \* \* /tmp/idoctor/nmon\_monitor.sh 300 12

d. Press the escape key to go back to command mode. If you made a mistake and need to fix it, use the applicable command from the cheat sheet referenced above to edit what was previously typed.

e. To exit and save changes, from command mode, type :x and press enter.

f. To exit without saving changes, from command mode, type :q! and press enter.

- g. Repeat the steps for all of the entries needed for NMON, NPIV, and SEA monitors.
6. Now that the monitors are configured, after each NMON, NPIV, or SEA collection the data will be copied from the collection directory to an SCPOUT directory. This will happen approximately every 60 minutes around the top of the hour. The SCPOUT directories are monitored every minute for new data. As new data arrives, it is sent to the specified IBM i remote directory and then deleted from the VIOS. Doing an 'ls' on the /tmp/vios\_investigator directory will show the applicable sub-directories which are created after the first time the scripts are run. This will happen at the top of the first hour after the entries are added to the Cron scheduler.

/tmp/vios\_investigator/nmon: NMON Monitor data is stored here  
/tmp/vios\_investigator/scpout\_nmon: NMON SCP output directory  
/tmp/vios\_investigator/npiv: NPIV Monitor data is stored here  
/tmp/vios\_investigator/scpout\_npiv: NPIV SCP output directory  
/tmp/vios\_investigator/sea: SEA Monitor data is stored here  
/tmp/vios\_investigator/scpout\_sea: SEA SCP output directory

#### **Section 4. Directory Monitor Setup Instructions**

1. Use the QIDRGUI/ADDMONDIR command to add the monitored directory or directories. Each entry will monitor the specified IFS directory for data arrival and import it into the specified library. Several examples are included below. Organize the data however you like.
2. Example: Entry to monitor directory for NMON, NPIV, and SEA files from VIOS1 and import to library VIOS1PERF. Each collection is imported and appended to existing files in the library such that there is one collection per day. As new data arrives from the VIOS throughout the day graphs can be refreshed to show the latest intervals.

QIDRGUI/ADDMONDIR DIR('/home/myuser/vios1') IFSPFX(VIOS1)  
 TGTLIB(VIOS1PERF) TGTPFX(VIOS1) IFSRTN(2) DBRTN(30)  
 MODE(\*DAILY)

```

Add NMON/NPIV Monitored Dir (ADDMONDIR)

Type choices, press Enter.

Directory to monitor . . . . . '/home/myuser/vios1'

File types to monitor . . . . . *ALL          *NMON, *NPIV, *SEA, *ALL
Prefix of IFS files . . . . . vios1         Character value
Target library . . . . . vios1perf        Character value
Target prefix . . . . . vios1            Character value
Days to retain IFS files . . . . . 2         0-9999
Days to retain DB files . . . . . 30        0-9999
Map files . . . . .

Import mode . . . . . *DAILY            *NORMAL, *DAILY
  
```

3. Parameter descriptions:

- a. Directory to monitor: IFS location to monitor for new files.
- b. File types to monitor: Type of VIOS data to monitor.
- c. Target library: Library where the data will be imported for analysis.
- d. Target prefix: This will be the prefix for the DB file members.
- e. Days to retain IFS files: How long to keep the raw data in the IFS. This data is not used by iDoctor for analysis purposes.
- f. Days to retain DB files: How long to keep the imported DB files. This data is used by iDoctor for various reports and graphs.
- g. Map files: If disk mapping files have been included from any IBM i partitions they can be referenced here.
- h. Import mode: \*DAILY appends each collection received into a single member per day. \*NORMAL means that each collection received has its own member. For example, collections received hourly at the end of the day would have one member using \*DAILY versus 24 members using \*NORMAL.

4. Start the directory monitor with the current user or a different user. Put this command in your startup to ensure the monitor starts up after an IPL or maintenance.

QIDRGUI/STRDIRMON USER(\*CURRENT)

5. Other applicable commands: ENDDIRMON, RMVMONDIR
6. Once files have been imported they can be viewed by opening any iDoctor component, such as Collection Services Investigator or Job Watcher, and expanding General Functions > Power > NMON, NPIV, or SEA > Analyze > Expand the library and member.

