Data Collection Guide for IBM N series 7-mode

Table of Contents

1.0 - Required data for troubleshooting N series

1.1 PSI required!

- Problem description (not all can be applicable):
- Is it a new install?
- Is it the first time that the problem occurs?
- If so, what could have triggered it? (Any code upgrade, config change, environment change)
- Has it already worked and when did it stop if yes?
- What are the symptoms? (hang, application not working, ...)
- When the problem occurs? (Morning, day, night...)
- Which volume/LUN is impacted?
- Which protocol is impacted?
- Which users/servers are impacted? (NetBIOS names, IP addresses)
- Current status of machine Environment:
- Versions of ONTAP, Snap Managers, Snap Drive, Hosts (i.e. Solaris /Windows, x64 /x86, release number)?
- Which applications are running?
- Which protocols are in use?

Business impact:

- What kinds of Services are impacted?
  - * Ex: Exchange mail server down, 400 users no email access.
- How many users impacted?
  - * Ex: 350 users of DB cannot work today
- How customers business is impacted?
  - * Ex: car spare parts manufacturing: SAP system is down, cannot print stamps orders, so cannot ship any parts to customers, all shipments are in hold. Customer will have to pay 20% fine for all delayed shipments.

1.2 Base data collection:

- Machine type, S/N, sysid, DOT version, config,
- Autosupport (which include above)
- Where is the data: dashboard or Ecurep

USE ecurep PMH link for data upload!

1.3 Base data collection Metro Cluster:

- Base data collection as for normal system and switches data collection (for fabric metro cluster)

PLEASE provide information if switches are under IBM support contract!
* Example for Brocade Metrocluster switches:

a) collect supportsave of the Metrocluster switches
b) clear the statistics using CLI: “slotstatsclear” and “statsclear” and "portstatsclear"
c) wait for 2 hours and than collect new supportsaves

- specify ISL lengths (distance between switches in the fabric) , and connections:
- patch panels, DWDM’s (type, transparent or not), cabletype, etc...

1.4 Base data collection Gateway:

Base data collection as for normal systems plus :

- switches data collection (if applicable)
- SAN drawing
- Backend storage firmware levels, and base data collection for the storage

if gateway is not booting :

- capture all serial console commands and outputs to a file
- capture one boot sequence, then
- go to maintenance mode and collect outputs of:

```
     >"Disk show -v"
     >"Storage show disk -p"
     >"Storage show adapter"
     > "Storage show port"
     >"Storage show initiators"
     >"Version –b"
     >"fcadmin config"
     > "aggr status -v"
     > "aggr status -r"
```

2.0 – DC Auto Support (ASUP)

After you have forced the ASUP to generate, please use Ecurep upload to pmh, to send the data to support.

ASUP data is stored under /etc/log/autosupport directory on root volume
* (example: /vol/vol0). It will create following directory and files.

```
yyyymmddhhmm.files      : This is a directory, several files under this
                         directory
yyyymmddhhmm.x          : Text file, command output
yyyymmddhhmm            : time and date which ASUP gathers data
yyyy                    : Year, i.e. 2012
mmdd                    : Month, date,  i.e 0419
hhmm                    : Hour, Minute  i.e 1710
```

How to force autosupport?

Connect to the nSeries controllers and issue the following command [options autosupport.doit IBM]. To force a autosupport while in takeover mode, go to the partner prompt with the command [partner] and issue the .doit

Note: For Filers/Gateways updated to DataOntap 8.x.x the message log is not any more included in the asup folder.
Please upload the messages logs from /etc :
/etc/messages …messages.0…..messages.5

1. CIFS share from Nseries available, Windows Client can access root volume.
   - In Windows Explorer go to “Tools” and “Map Network Drive”
   - `map \\xx.xx.xx.xx\c$` (XX.XX.XX.XX ==> filer IP adresse)
   - Copy directory and files under etc\log\autosupport folder (yyyymmddhhmm) and send them to Ecurep

2. NFS export available, UNIX/Linux Client can mount root volume
   - Mount the system volume, typically **`/vol/vol0`** to any mounting directory (i.e. /mnt)
   - Copy directory and files under **`etc\log\autosupport`** folder (yyyymmddhhmm) and send them to Ecurep

3. No CIFS or NFS available

Using Nseries Console (Serial or telnet), issue following commands and capture output.

Set the options:

- Options autosupport.enable on
- Options autosupport.support.enable on.

Enable Console Log function (to capture output of screen to file) In general most of ASCII Emulator Program, like Hyper term, Teraterm or Putty have this capability. Please make sure you have the console connection at full screen mode.

```
>"priv set advanced"
>"options autosupport.doit IBM"
>"ls /etc/log/autosupport"
```

You should see file name, yyyymmddhhmm.x. See the explanation above for the naming convention.

```
>"rdfile /etc/log/autosupport/yyyymmddhhmm.0"
>"rdfile /etc/log/autosupport/yyyymmddhhmm.1"
>Issue same command as above, yyyymmddhhmm.2, yyyymmddhhmm.3,
  and so on
```

Stop Console log, copy and paste the “rdfile” contents to single files and send them to Ecurep.

Please use ecurep to upload the data to IBM. See section in the Cookbook for the procedure

4. Using FTP for asup collection

On the filer set:

```
 "options ftpd on"
```

Enable anonymous FTP on the filer:

```
>"options ftpd.anonymous.enable on"
>"options ftpd.anonymous.home_dir /vol/<root volname>"
>"options ftpd.anonymous.name anonymous"
>"options ftpd.enable on"
```

Create the /etc/passwd file for anonymous FTP.

```
>"priv set advanced"
>"rdfile /etc/passwd"
>copy the content to a text file and add a line with
"ftp::65533:65533:FTP Anonymous:/home/ftp:"
>filer*> "wrfile /etc/passwd"
>copy the content of above textfile and press enter, before
leaving with "Ctrl C"
"ftp::65533:65533:FTP Anonymous:/home/ftp:"   (presse enter,
and "Ctrl C")
>check with "rdfile /etc/passwd"
```

FTP access to root volume is now possible for user: 'anonymous'

3.0 - DC for performance (perfstat)

-----PERFORMANCE CASE TEMPLATE -----

1a) what performance-affecting behaviour are they seeing? (Latency values, etc in seconds or ms. if errors, then the exact errors seen, and where)

1b) what is this performance-affecting behaviour causing problems with?
(Applications, resources)

2) What behaviour are they expecting to see? How fast is it on normal days, etc.
(What is slow? Are reads slow / writes slow / browsing of file or directories /file creations/mounts? How are you judging "slow"?)

3) When did it start?
(Was performance acceptable at one time? / New install and/or performance was never good / degradation over time?)

4) When does it happen?
All the time, only at certain (production) times, etc.

5) Who is it affecting?
All users?
Users of a particular CIFS-based application? Etc

6a) what has changed between when it was working fine and now?

6b) Have any new errors or changes been noticed in the filer's autosupports?

7a) any errors appearing on the filer?

7b) any errors appearing on the clients?

8) What kind of client machines are working with this filer?
Is it straight cifs, cifs/nfs, cifs/SAN, only SAN (to a cifs/win machine), etc?

-----PERFORMANCE CASE TEMPLATE -----

Note: a perfstat my only be generated / uploaded for analysis upon request by level 2!

Performance and Statistics Collector (perfstat) - 7.39 UNIX/Linux, 7.39 Windows
- This Perfstat version is supported for DataOntap 7.x and 8.x 7-mode Code Versions

Note: Please always verify that you are using the most recent version of Perfstat. Older versions may trigger bugs which degrade filer performance or stability.

http://support.netapp.com/NOW/download/tools/perfstat/ (*)
(*) The tool will be downloaded and provided to customer by IBM Level 1 support.

For RSH setup on the filer, please see KB 1010082:
https://kb.netapp.com/support/index?page=content&id=1010082&locale=en_US

Perfstat is a data collection tool with several key properties:
- Captures all needed performance information with one command
- Captures information from host(s) and filer(s)
- Captures all information simultaneously for cross correlation
- Operates on all host platforms and all filer platforms
- All captured data is recorded in a single plain-text output file

Perfstat comes in exactly two flavors:

1. Unix/Linux version (perfstat.sh)
2. Windows version (perfstat.exe)

Supported platforms:

- AIX, ESX, HP-UX, Linux, OSF1, Solaris, FreeBSD
- Windows: 2000/XP/2003/2008*/7*

Note: Windows 2008 and Windows 7 support requires openSSHforWindows to be setup. Please follow the procedure available in "openSSH_procedure.txt" file bundled with perfstat7 windows binary to achieve the same.

Here is one typical command:

```
>"perfstat -f <filer_IP> -l root:password -t 4 -i 5
>perfstat_filer_name.out"
```

Note: For Windows users the example below shows a nice way to format perfstat output file names with date and time information:

perfstat -f filer1 -t 3 -i 5 > perfstat-"Date_%date:~4,2%%date:~7,2%%date:~10,4%-Time_%time:~0,2%h%time:~3,2%m%time:~6,2%s"

4.0 - DC for Snapdrive and Snapmanager Problems

4.1 - Snapdrive for UNIX / Windows

To collect snapdrive data for UNIX and Windows, provide the customer with the 'snapdrive.dc'(Unix) or the 'ONTAPWinDC' (Windows) tool:

Unix tool:
http://support.netapp.com/NOW/download/tools/snapdrive_data_collector_unix/ (*)

Windows tool:
http://support.netapp.com/NOW/download/tools/snapmanager_e2k_dct/ (*)

(*) The tool will be downloaded and provided to customer by IBM Level 1 support

4.2 - Snapmanager for Exchange, SME

Beside the output of the ‘ONTAPWinDC’ we should collect the logs from SME, here the ‘Report’ folder should be zipped and provided:

```
'C:\Program Files\IBM\SnapManager for Exchange\Report'
```

4.3 - Snapmanager for SQL, SMSQL

Beside the output of the ‘ONTAPWinDC’ we should collect the logs from SMSQL, here the ‘Report’ folder should be zipped and provided:

```
'C:\Program Files\IBM\SnapManager for SQL Server\Report'
```

4.4 - Snapmanager for Oracle, SMO

Windows version

```
'C:\Program Files\Ontap\SnapManager for Oracle\bin'
'C:\Documents and Settings\Administrator\Application
Data\Ontap\smo\<smo version>'
```

Unix Version
collect a data collection depending on the Unix OS system, on request.

Depending on the OS we also need either the **‘ONTAPWinDC’ or the ‘snapdrive.dc’ log.**

Data collection for SMO is based on ‘dump’ cmds against the Operation, the Profile and the System.

```
- "smo operation dump -id guid"
  -> "smo_dump_operation-id.jar"

- "smo profile dump -profile profile_name"
  -> "smo_dump_profile-name_host-name.jar"

- "smo system dump"
  -> "smo_dump_host-name.jar"
```

Locate dump files (SMO 3.0.0 example)

If using the graphical user interface:

```
'user_home/Application Data/Ontap/smo/3.0.0/'
```

If using the command line interface:

```
'user_home/.ontap/smo/3.0.0/'
```

Add log files

SnapManager records all log entries and places them into one set of rotating log files.

The log files are found in:

```
-> '/var/log/smo'
```

4.5 - Snapmanager for SAP, SMSAP

Depending on the OS we need also either the **`‘ONTAPWinDC’ or the ‘snapdrive.dc’ log.`**

Data collection for SMSAP is based on ‘dump’ cmds against the Operation, the Profile and the System.

```
- "smsap operation dump -id guid"
  -> "smsap_dump_operation-id.jar"

- "smsap profile dump -profile profile_name"
  -> "smsap_dump_profile-name_host-name.jar"

- "smsap system dump"
  -> "smsap_dump_host-name.jar"
```

* Locate dump files (SMSAP 3.0.1 example)

If using the graphical user interface:

`'user_home/Application Data/Ontap/smo/3.0.1/'`

If using the command line interface:

`'user_home/.ontap/smo/3.0.0/'`

Add log files

SnapManager records all log entries and places them into one set of rotating log files.

The log files are found in:

`-> '/var/log/smsap'`

4.6 - Snapmanager for VI (SMVI) and VMware ESX logs

https://support.netapp.com/Knowledgebase/solutionarea.asp?id=kb46059 (*)

(*) The tool will be downloaded and provided to customer by IBM Level 1 support

To collect SMVI logs:

On the Server where SMVI is installed, collect all of the logs in the following location.

```
  'C:\Program Files\IBM\SMVI\server\log'
```

To collect VMware ESX logs:

Option 1: Upload the logs either using the vm-support utility:

1. Log into the service console with root access
2. Type "vm-support" without any options.

The utility will run and create a single Tar file that will be named

`'esx---..tgz'.`

Option 2: Generate the same file by using the VMware Infrastructure Client

(VI Client).

1. Select Administration
2. Select Export Diagnostic Data

4.7 - Snapmanager for Microsoft Office SharePoint Server, SMMOSS

SMMOSS data collection is in most cases 'GUI driven'.

On the SnapManager for Microsoft Office SharePoint Server GUI you will find on the left-hand the different 'control-task' which you need to select also for data collection.

The SMMOSS does provide the logs on log levels.

If cus can recreate the issue, please change the log level.

-----SMMOSS CASE TEMPLATE -----

There are two server managers needed for SMMOSS

- Servermanager
- Provide Scheduled backup at the DB Item-level
- Real time restores
- Automated data pruning
- Central administrator of multiple sharepoint sites
- SharePoint Server Agent
- SharePoint Control Agent
- SharePoint Member Agent

1. What part of SMMOSS is having the problem? (See above for components)
2. Is this a new Installation or has this been working before?
3. Where there any changed in the system, application, network or other in their environment?
4. What is the problem
   a. Backup
   b. Restore
   c. Access
   d. Communication
   e. Browser issues
5. What (if any) are the errors messages seen?
6. How many projects / objects are in the SharePoint Database?
7. When was the last successful backup and restore?

-----SMMOSS CASE TEMPLATE -----

Data collection needed to analyze!

1. Generate a new AutoSupport (See section in the cookbook)

a) Changing Log Levels

Change Manager component’s log levels via Control Panel, Log Manager

Change Agent log level via Control Panel, Control service, Agent Monitor

Logs from the SMMOSS Job Details:

Select **'Job monitor'** on the left -> View, there will be some zip files containing the SMSQL CLI commands used and the CLI output

b) Downloading SMMOSS Logs

Under Control Panel, Log Manager

Select the agents or the services on the right

"Download" the logs to get a zip file of the various logs

c) SMMOSS Interface Issues

Problem: The SMMOSS interface does not launch properly what log files do I retrieve?

File name: 'NetApp-SMMOSS.log'

Location:

The file is located on the machine with the SMMOSS Manager installation

By Default: **'../Program Files/IBM/SnapManager for SharePoint Server/VaultServer/WASCE/logs'**

Problem: The SMMOSS reports are not loading properly what log files do I retrieve?

File name: 'NetApp-SMMOSS.log'

Location:

The file is located on the machine with the SMMOSS Manager installation

By default: **'../Program Files/IBM/SnapManager for SharePoint Server/VaultServer/WASCE/logs'**

Note: Toggle the **'NetApp-SMMOSS.log'** logging to debug through the SMMOSS interface

d) SMMOSS Backup Issue

Problem: If a backup fails to generate the index file what SMMOSS log files do I retrieve? File name: `'NetApp-SMMOSSmedia.log'`

Location:

The file is located on the SMMOSS Media Server installation

By default: `'..\Program Files\IBM\SnapManager for SharePoint Server\VaultServer\SMMOSSMedia\logs'`

Note: Please toggle the 'SMMOSS-media.log' logging to debug through the SMMOSS interface

Problem: If a backup fails to generate the index file what SMMOSS log files do I retrieve?File name: `'SMMOSS.evt'` from the Event Viewer

Location:

The file is located on the front-end web server Use Event Viewer

Look for a NetApp event – save it as an `'.EVT'` extension.

Note: Please toggle the SMMOSS.evt logging to debug through the NetApp interface

e) SMMOSS Restore Issue

Problem: If a restore fails what SMMOSS log files do I look at for error message?

Location:

Retrieve the SMSQL and SnapDrive output log which is attached to the restore job's

`'Job Details'`, see above. Look for the command that is executed during the restore process

Problem: If a restore fails what SMMOSS log files would have the error?

File name: `'SMMOSS.evt'` from the Event Viewer

Location:

The file is located on the front-end web server Use Event Viewer

Look for a SMMOSS event and save it as an `'.EVT'` extension

Note: Please toggle the logging to debug through the SMMOSS interface, see change log levels above

Problem: Customer attempts to restore lists or document libraries using only the

NOT OVERWRITE option, but restore job fails to restore meta data from the list or document library

Resolution:

The NOT OVERWRITE option does not provide a strong enough option for the Restore Job

If the restore using the NOT OVERWRITE fails, attempt the Restore Job again using the OVERWRITE option

The OVERWRITE option will ensure that all meta data is restored properly

5.0 - DC for NDMP troubleshooting

Base data needed:

1. Auto Support from the filer in question
2. SAN / Ethernet layout
3. Description of the NDMP process from the start until the storage
4. ndmp debug log
5. Dump to Null test (see description below)
    - Console LOG during the Dump
    - The LOG Backup file

NDMP debug log:

Set the NDMP debug level to 70 on the filer, and perform the NDMP operations

- **"ndmpd debug 70"**

Once the problem has been replicated, turn off NDMP debugging using the command:

- **`"ndmpd debug 0"`**

Note: The following tests to null can be run with the debugging enabled to help troubleshoot the performance problem.

To test the native read performance from the storage system, initiate a level 0 dump to null of the affected volume.

This will test the storage system’s performance reading the selected data from the disks and help isolate whether

the problem is with the disk reads.

Note that this should be done during a period of time when the workload on the storage system is equivalent to the workload during the problematic backup.

Connect to the filer and enable Console Log function (to capture output of screen to file) in general most of ASCII Emulator Program, like hyper term, Teraterm or Putty has this capability. Issue the following command

```
>"rsh storage_system_IP dump 0f null /vol/<volname>"
```

Note: Depending on your configuration, dumping to null can put a great load on your system, which may lead to an impaired end-user experience. It can, however, be aborted anytime (By using the "CTRL+C" key combination). If the command is issued from the command-line interface (instead of an 'rsh' command, for example) it will block the console until finished.

Collect the NDMP and backup logs from the storage system.

- **`'/etc/log/ndmpdlog.date'`**
- **`'/etc/log/backup'`**
- Generate a new AutoSupport:
  `"options autosupport.doit ndmp_test"`
- Collect the logs from the backup applications for the same period as the storage system logs.

Please use ecurep to upload the data to IBM. See section in the Cookbook for the procedure

6.0 - DC for DFM/OM and OnCommand Unified Manager (core)

Please download the tool "**`DataFabric® Manager Data Collector (dfmdc)`**`3.0`"

(Or the latest available) under following address and use the option to generate the data form the system.

https://support.netapp.com/NOW/download/tools/dfmdc/ (*)

(*) The tool will be downloaded and provided to customer by IBM Level 1 support

Overview

DataFabric® Manager Data Collector (dfmdc) is a tool to collect data used by IBM support when troubleshooting DataFabric Manager issues. The tool runs various dfm command, gathers up the log files, then consolidates them into a single compressed tar image which can be uploaded to IBM for analysis. DataFabric Manager license is required on DFM host.

Installation

Simply **`'unzip'`** (untar) the distribution somewhere on the DFM host. No other configuration is needed. Included in the distribution includes Perl source code as well as binaries for both Windows and Solaris for those systems who may have difficulty running Perl scripts.

Usage is as follows:

```
'dfmdc[.pl] [-nlq] [-d dir] [-c case#] [-f filer[,filer...]]'
```

The .pl extension is needed if using the Perl source. It is not needed if running the binaries.

`'-n' :`

This flag does not gather the DFM logs. By default, the tool gathers the logs.

`'-l'` :

This flag *only* gathers the DFM logs and the command outputs will not be gathered.

`'-q'`:

Quiet mode. By default, the tool will print out progress. If this flag is set, no such output will be displayed

`'-d'` dir :

This flag allows the user to specify an output directory for the output file, as well as the temporary files which are deleted when the tool completes. By default, it uses the current working directory.

`'-c'` case #:

This flag will prepend the given case # to the name of the final tarball that is created. This is useful if the tarball is to be uploaded to IBM for analysis.

`'-f'` filer: By default no filer diagnostics are gathered. This flag will collect a dfm host diag for each filer listed.

7.0 - DC for Packet trace procedure

In case the customer do report **`‘IP related’`** problems like general connectivity issues, access problems via CIFS/NFS, it is sometimes useful to have a ‘LAN’ trace from the Filer.

The filer does provide a trace facility to capture such trace, the cmd to use is the

`‘pktt’` cmd

The cmd itself is described in more details in the **`‘cmd-reference guide, pg.389`**:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S7002461&aid=1

The sysntax for to start a `‘pktt’` trace is:

```
"pktt start {if | all} [-b bsize] [-d dir] [-s size] [-m pklen]
[-v] [-i ipaddr] [-i ipaddr] ..."
```

In most of the cases it is absolutely sufficient enough to leave the default values for the option parameters.

* Example:

To trace the IP traffic on the vif interface **`‘vif1’`** towards a client with IP addr. 10.10.10.1 and to save the file to directory **`‘/etc/tmp’`** with the size of 200mb you would start the trace:

```
"pktt start vif1 -d /etc/tmp -s 200m -i 10.10.10.1"
```

To stop the trace simply type:

```
"pktt stop {if | all}"
```

When tracing to file, you will find the trace in the directory you specified by the ‘pktt start’ cmd, the naming will be:

```
‘ifname_yyyymmdd_hhmmss.trc’
```

In our example it would be -> **`'vif1_20091006_162637.trc'`**

Add info:

To trace on the client site at the same time as on filer site you can use for ex. the tool ‘**`wireshark’`** which is available for several OS, see:

http://www.wireshark.org/

Please use ecurep to upload the data to IBM. See section in the Cookbook for the procedure

8.0 - DC for Snapmirror

Data required:

1. Detailed description on the snapmirror issue
   - Description about problem snapmirror volumes (source and destination)
   - did the relation work before?
   - What was changed?

2. Auto support for source and destination filer
   - In addition the '/etc/hosts' and '/etc/rc' files for both filers

3. on destination filer

```
-  '/etc/snapmirror.conf' file
-  '/etc/snapmirror'  log file
```

4. on source filer

```
- '/etc/snapmirror.allow' log
```

9.0 - DC for Snapvault

Data required:

1. Detailed description on the snapvault issue
   - Standard Snapvault or Open Systems Snapvault (OSSV)?
   - Description about problem snapvault volumes/qtrees (primary/secondary)
   - did the relation work before?
   - What was changed?

2. Auto support for source and destination filer and in addition

- `'/etc/hosts'`
- `'/etc/rc'`
- `'/etc/export'`
- `'/etc/log/ems'`
- `'/etc/snapmirror.allow' or 'snapvault.access'`
- `'/etc/log/snapmirror' logs file on the destination.`

3. For OSSV information about OSSV version and Server Version needed.

- start OSSVINFO.exe (Windows) or OSSVINFO.pl (unix) and collect the OSSV info data , if this is not working collect the OSSV Server log in

```
'Windows: C:\Program Files\IBM\snapvault\etc<'
'Unix/Linux: /'
```

10.0 - DC Core file

10.1 - how to upload

The core files are saved in the /etc/crash directory on the filer. Depending on the OS version the core files saved are named 'core.x' and either `'core.x-small' or 'mini-core.x'`,

where x is a number. If you have a `'core.x-small' file, or 'mini-core.x'` file it is generally not adequate for analysis, so it is a safe practice to upload the larger core file.

The files in the folder are only for diagnostic purposes and can be safely deleted after the problem has been resolved. The steps to upload a core file are outlined below.

Note: Verify that the core file is valid! (See the next section for that)

Run the `"savecore -l"` (check mode) command on the filer. If the proper panic string appears in the output, then the core-header is intact and can be read by NetApp (although it could still be incomplete).

In the absence of core-file, get the following:

- the `'EMS'` and `'Audit'` logs, from the `"/etc/log"` folder.

Prepare the core file for upload

On a CIFS or NFS licensed filer, complete the following steps to prepare a core file for upload and analysis:

1. For filers, from the CIFS or NFS admin host, navigate to the `/etc/crash` directory. Locate the most recent corefile using its timestamp and save it to the admin host. Rename the core file so that it contains the PMR number,

using the following naming convention:

`'[PMR_number].core.[Sequence_number].nz'`

* For example, if the core file name is "core.2.nz" and the case number is

`'1234.567.890', rename the 'core.2.nz' file to '1234.567.890.core.2.nz'.`

> Note: Any other file that you need to upload can be renamed in this fashion. Appending the PMR number to the file name helps speed the identification of the file and expedite analysis process

If CIFS or NFS is not installed on the filer, use either of the following options to get the file onto another local machine:

OPTION 1:

Follow these steps to to easily download core using the httpd admin front end.

1. Set `"options httpd.admin.enable true"`.
2. If access to particular hosts needs to be limited, set `"options httpd.admin.access"` to

   `'host=<host1>,<host2>,...'`
3. Go to `'/na_admin/cores>http://<system>/na_admin/cores' and download the 'core'.`

OPTION 2:

Using the ftpd admin front end, you can easily download any cores and minicores you need.

1. Enable FTP access on the filer **`"options ftpd.enable on"`**
2. Set the default home directory to '/etc'
   `"options ftpd.dir.override /vol/<root volname>/etc/crash"`
3. Connect to the filer using FTP from a local client.
   Before invoking the FTP command, be sure to note the directory you are in.
   This is the where the core file will be when you GET it.

4. Issue the LS command to find the necessary file.
   * For this example the core file is named **`'core.2.nz'`**.

5. Issue the GET command to upload the core file to your local computer.

```
>"get core.2.nz"
```

6. Terminate the FTP session.

7. Turn off FTP access on the filer.

```
>"options ftpd.enable off"
```

8. Make sure that you are in the directory that the file is in,
   then rename the file using the naming convention specified above.

Note: Do not upload the 'mini-core' file unless specifically requested to do so by IBM support.

Please use ecurep to upload the data to IBM. See section in the Cookbook for the procedure

10.2 - split, reassemble and verify

A core file is typically around 20 percent of the storage box memory; sometimes a timeout occurs when uploading a core file. One solution is to split the corefile into smaller pieces. The best size for each piece depends on the bandwidth of the upload link. Splitting the corefile into pieces that are 1 GB in size is recommended, but the pieces may need to be smaller if the upload link is very slow.

This procedure consist in XXX steps

a) Splitting a core file on a UNIX host

* Here is an example of how to split a core file named core.0.nz for case # 12345 into pieces that are 1GB each:

client-unix@13:26:24{138} **`"split -b 1024m core.0.nz 12345.core.0.nz".`**

List the files and note the order in which the split files appear. Later, you'll need to re-assemble them in this order, using the "cat" command.

client-unix@13:30:25{139} **`"ls -l"`**

total 8633168

```
'-rw-rw-r-- 1 owner engr 1073741824 Dec 7 13:27 12345.core.0.nz.aa'
'-rw-rw-r-- 1 owner engr 1073741824 Dec 7 13:28 12345.core.0.nz.ab'
'-rw-rw-r-- 1 owner engr 1073741824 Dec 7 13:29 12345.core.0.nz.ac'
'-rw-rw-r-- 1 owner engr 1073741824 Dec 7 13:30 12345.core.0.nz.ad'
'-rw-rw-r-- 1 owner engr 116523035 Dec 7 13:30 12345.core.0.nz.ae'
```

Once the process has completed, all pieces can be uploaded in parallel if the upload link has enough bandwidth.
For low bandwidth upload links, each piece can be uploaded separately.

b) Splitting a core file on a Windows host

On Windows systems, there are various Shareware programs that can split files. The following example shows how to use to use the popular program WinZip:

1. **`"Zip core.0.nz" into 'core.zip'.`**
2. open **`'core.zip' with 'WinZip'.`**

3. Open the Actions menu and choose "Split". Specify a filename '12345core' and the desired size for each piece (1 GB). Click `"OK"`.

Upload the pieces and let the case owner know how many there were.
Verifying the core file.
After the core file has been reassembled, verify that it is a usable core.

Under UNIX:

Two options are presented:

- The first requires the customer to run a command and supply the output.
- On the client machine, enter the following command:

```
client-unix@13:30:25{139} "md5sum core.0.nz"

'e5ed9256124688bd3860862971ff9c6f core.0.nz' <----
```

On the NetApp support UNIX host:

```
netapp-unix@13:30:25{139} "md5sum core.0.nz"

'e5ed9256124688bd3860862971ff9c6f core.0.nz' <----
```

Compare the output from the two sources; if it is the same on the customer machine and the NetApp support host, the corefile has been reassembled correctly.

- The 'coretool' option does not require additional customer info.
  If the "coretool" utility says it is corrupted (unreadable),
  we can conclude that it is a useless core.

```
netapp-unix@13:30:25{139} "coretool core.0.nz"
```

Under Windows:

Download 'md5sum.exe' and save it to the directory where the core file is located. Alternative download the 'fciv.exe' from Microsoft and follow the described procedure.

http://support.microsoft.com/kb/841290

- At the command prompt run 'md5sum.exe' on the core file to be uploaded.

* Example:

```
"C:\md5sum -b core.0.nz"
```

The output will look similar to:

```
'e5ed9256124688bd3860862971ff9c6f *core.0.nz'
```

Copy the 'md5sum.exe' output and send to E for comparison.

- Upload the core file to Ecurep for analysis using the directions provided in the cookbook

IMPORTANT:

Windows uses a BINARY MD5 checksum calculation by default as indicated by the "*" in the output. Use the "-b" option under UNIX or Linux to do a binary calculation. If the output doesn't show '*core.0.nz' then the calculated checksum is a TEXT, not BINARY calculation.

Note: You must be in the directory where the file resides. For more information, see How to find a core file.

- On Jelly, after upload, run md5sum.exe against 'core.*.nz'.

* Example:

'[user@jelly]$ md5sum -b core.0.nz'

The output will look similar to: '*core.*.nz'

11.0 – DC System Manager (OnCommand)

It is not reported anywhere on 'OnCommand System Manager' Release Notes and Quick Start Guide, that is possible to generate a Support Bundle.

To get it, select Help button on top left; choose option Support Bundle and click on:

- **`'Click here to generate the file'.`**

A ZIP file is obtained with some files.
They could be useful to troubleshoot access problem on system.

12.0 - Issues with connected Tapes

In addition to autosupports, provide this filer outputs:

```
>filer> "storage show tape supported -v"
>filer> "sysconfig -m"
>filer> "sysconfig -t"
>filer> "storage alias"
```

- If the tape drive is fabric connected , please provide logical diagram of fabric and zoning.
- If the tape drives are not seen, log into the switch and gather the switch information and send it:
  - **`BROCADE "supportshow"`**
  - CISCO **`"show tech-support"`** details
  - MCDATA **`"show all nsshow"`**

Send the output to ECuRep .
Provide also the complete SAN Layout.

13.0 - nSANnity data collection tool

Please provide the nSANity output on support request only:
https://support.netapp.com/NOW/download/tools/nsanity/ (*)

(*) The tool will be downloaded and provided to customer by IBM Level 1 support.

nSANity Data Collector is a support tool designed to aide users and technical support in troubleshooting complex issues.

nSANity is able to collect diagnostic and configuration data from a variety of components including:

Data ONTAP Storage Controllers
Windows 2003 and 2008 hosts
VMware ESX hosts (excluding i variants)
Linux hosts with kernel 2.6
Solaris hosts
AIX hosts
HP-UX 11i hosts
Brocade switches
McData switches (EOS, EOSn)
Cisco switches (IOS, NXOS and SANOS)
QLogic switches

nSANity is designed to perform data collection from the components remotely using a variety of network protocols and native authentication mechanism. The nSANity executable binaries are produced for Windows, Mac OSX and Linux.

The Mac OSX and Linux versions of nSANity cannot communicate with Windows components, though the Windows version of nSANity can communicate with all component types with out exception.

When remote connectivity is not available, nSANity may be run directly on the target component if that component is Windows or Linux.

nSANity supports the following network protocols and will prefer the secure protocol when available.

`'HTTP/HTTPS'` is used for communicating with Data ONTAP

`'SSH'` is used for communicating with Cisco, Brocade, VMware ESX, Linux, Solaris, AIX and HP-UX.

'WMI' is used for communicating with Windows hosts

A telnet fallback is available for McData switches when 'SSH' is not enabled

* Example on nSANity command:

* Example Usage:

To collect SnapDrive debug data and all general data from a Windows host called 'myserver' and you are a user with administrative rights use the following syntax.

```
-     'c:\> nsanity.exe -S snapdrive windows://myserver'
```

Or to collect SnapManager diagnostic data which will include SnapDrive data.

```
-    'c:\> nsanity.exe -S snapmanager windows://myserver'
```

If the Windows host is part of a cluster then add all cluster nodes.

```
-    'c:\> nsanity.exe -S snapmanager windows://node1
windows://node2 windows://node3'
```

Components may be collected at the same time or separately, either way will produce the desired results.

* Example for nSanity output, included are Filer, ESX, and Metrocluster Switches and Switches to the Host

```
'c:\> nsanity.exe -S general ontap://root:*@<filer1>
      ontap://root:*@<filer2>
      brocade://admin:*@<MCswitch1>
      brocade://admin:*@<MCswitch2>
      brocade://admin:*@<MCswitch3>
      brocade://admin:*@<MCswitch4>
      vmware://[username[:password |*]@]server[:port] .......'
```

Note: Upload all data to the IBM EcUREP Server.
Please read the file "**Data Transfer to IBM.pdf**" for more details.