whNvInfo

Network NVIDIA GPU Monitoring

Documentation

Version 12

Copyright © 2015-2018 by Wolfgang Hurst

https://whurst.net

Introduction

What is whNvInfo

whNvInfo is a small program that provides several environment sensor data on screen. It just for NVIDIA Graphic Cards (do not ask for AMD/ATI or Intel – I don't own these).

whNvInfo can display and collect data from the network too.

The Visual Display

The visual screen separated by "Panels" or "Fields" or however you will call it. The first Panel is System related, the others are for NVIDIA GPU cards, if any.

Here a sample how it could looks like:

```
SYSTEM              GeForce GTX 780 Ti   GeForce GTX 780 Ti   GeForce GTX 680
   CPU Load%    92      GPU Load%    0       GPU Load%    0       GPU Load%    0
   CPU Mhz    3217      Core MHz   324       Core MHz   324       Core MHz   324
     Procs      87      Mem MHz    324       Mem MHz    324       Mem MHz    324
   Threads    1202      Free MB   2543       Free MB   2752       Free MB   4027
P Free MB    40652         Fan %     0          Fan %     0          Fan %     0
V Free MB    72150      Core °C     41       Core °C     43       Core °C     44
Net IO Mb        0      Mem °C      41       Mem °C      43       Mem °C      44
```

You see the System values on the left side, and 3 GPUs. Network servers displayed in the next row and so on.

History

The versions up to 6 are not network enabled versions. They are still available on my website. Please take a look to there to get the features, they didn't describe here.

Version 7 is the first version with network functionality, and the first version with this documentation document. For more details, please take a look to the website, where describes all features.

Upcoming versions will be detailed describe here at the end of the document.

System Panel Values

CPU Load%

Display the current system CPU load. This value is an average across all cores.

CPU MHz

Display the current average MHz across all cores. Meanwhile the CPU will fall down in the frequency if nothing to do.

Procs

Display the count of all processes in this system.

Threads

Display the count of all threads in this system.

P Free MB

Display the maximum physical free memory. The display will change to xxx G if you have more than 99999 MB physical memory free.

V Free MB

Display the maximum virtual free memory. The display will change to xxx G if you have more than 99999 MB virtual memory free.

Net IO Mb

Display the current network throughput in Megabit per second. Across all network interfaces.

NVIDIA GPU Panel Values

GPU Load%

Display the current GPU CPU load. This value is an average across all cores.

Core MHz

Display the current core MHz of the GPU.

Mem MHz

Display the current memory MHz of the GPU card memory.

Free MB

Display the maximal free memory available for display and stuff. For TCC mode this will still be zero, because TCC has no video memory.

Fan % or Fan RPM

Display the current fan speed. I didn't have any fans anymore, on my examples I just have 0 % fan speed, because of water cooling.

Core °C

Display the core temperature in Celsius.

Mem °C

Display the memory temperature in Celsius. Theoretically Core and Mem should very close, but if you have a serious problem than it will have a big gap. Than you should take a look into the case to see what happen.

Panel Color Styles

Values

All values they have a scale, like memory, MHz, or load are internally calculated to percent. Each percent group has another color.

Starting with gray and green, that means this is idle, or do nothing.

Orange is over 60% and red will start at 95%.

But in most cases you simply have red or green.

Network Colors

If you have network server configured, the panels become other colors to signal an error.

Normal working and enabled server are display normal, like the system internal panels:

Deactivated servers displayed darker, like this:

Servers with any kind of problems are red, like this:

You can turn off the display of failed and / or deactivated servers in the network setup.

Layout

The layout can changed from horizontal to vertical display.

Horizontal display means each system or server start with the system panel on the left and all GPU panels added to the right. Like this:
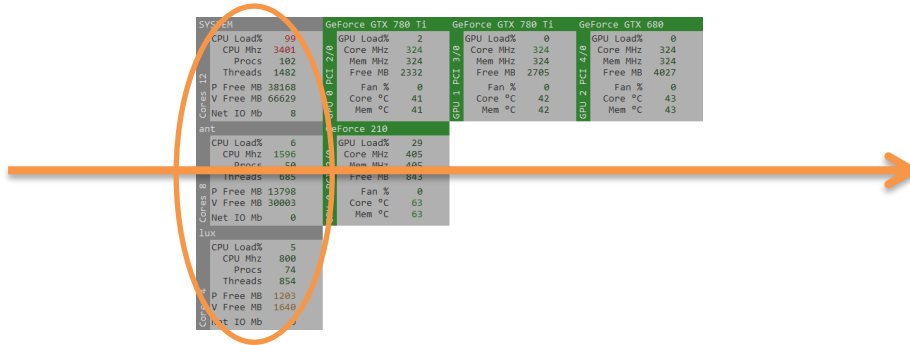
Vertical display means each system or server start with the system panel on the top and the GPU panels will add on the bottom, like this:
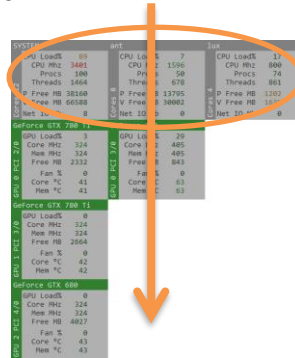
You can change the layout by dragging the display to the boundaries of the screen, of the full virtual screen. Drag it on the left or right side to become the vertical layout, drag it to the top or to the bottom to get the horizontal layout.

If you have a multi monitor setup with different screen sizes and monitors, like this:

You will have problem to reach the bottom on screen 2 in this example. I didn't catch this. Use the top in this example. If you ever wonder why it is not working, now, you know why.

Setup

General Setup

The general setup contains the font size and window placement setup.

Font size Setup

You can use the slider or enter the font size into this field. The display is immediate updated. It uses the Windows System Font Consolas. If you start the program the first time ever, it uses 12px as default.
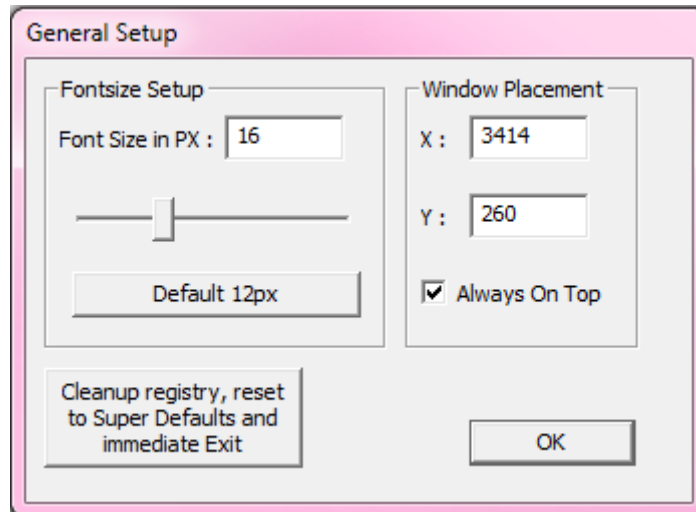
Window Placement

The X and Y fields show you the current coordinated of the window. You can enter your own coordinates if you like, or you can drag the display while press the mouse button. These two fields may help you to bring a display back to your screen. Sometimes on multi monitor setups not all the time all monitors are running, or you run it on a temporary display and you disconnect it. Here the change to bring the window back to screen.

Always On Top

Check this to make this window as a top level window. This causes sometimes problems with video player, they are super top level and sometimes the old top level windows are gone. If you have this situation, open this setup and uncheck and check back again this setting to bring it back to front.

Cleanup registry reset and exit button

Press this and the registry is gone and the program exists. It is like uninstall, but that one works and really cleans up itself.

Network Setup

Introduction

The whNvInfo can provide all monitoring values per network. It uses the UDP protocol. The server is not start per default and must explicit activated. Once activated it run automatically on start. This setting is stored in the registry.

There are two settings, one you can setup the server and the other one is for setup the client. Both setups combined in one dialog. Do not getting confused. The top is for the server, the bottom for the client.

Server Setup

Picture 1 - Network Server Setup Elements

Enable Network Functionally

This checkbox has two functions. If it checked, it gray out the values hostname and port and start up the server. If it unchecked, the server will stop and the fields getting editable back again. That means two functions with just one check box.

Visible Hostname

This is the hostname given to the clients. You can choose any name you want. This hostname is only given to the clients they ask. It has no other impact.

UDP Port

This is the port number for the server to use. You should use a port that not in use of cause. An error message pops up if the port already in use, but the message is sometimes a little confused. You can use the ports from 1024 up to 65535.

Server Status

The server status field informs you about the status of the server. It can be "Running" or "Stopped" or "Error" or other. The Server Status will be reflected on screen with different colors. Please take a look into the Display Chapter.

Server Uptime

This field is filled if the server running, where you can see how long the server runs without crash.

Server Queries

There is an internal counter that counts each valid query from the clients. This value displayed there. It could be useful if something not working, there you can see if queries arrive or not.

Show Server Detail Status

A message box with detailed information will pop up. This is useful if the server status set to "Error" or something similar.

Client Setup

Introduction

The client setup is much more complicated as the server setup. You can add a server, remove it and move it in the list, can setup some special values on it. It looks complicated, but I try to help you by enabling and disabling all controls they make no sense in all situations. I hope this will be helpful. At the beginning the list is empty and the most controls are disabled. You first need to add a server. You can use an IPv4 or IPv6 address, or the hostname with domain, or not. It follows the regular name resolving on the command line. Please notice IPv6 will be preferred if the host have a valid IPv6 address, there is no fallback to IPv4 at this moment.

You can add 64 servers. This is an internal hardcoded limit. If you have more servers, and a really big screen, ask me, I can modify it, but I guess 64 servers are enough.

For setup the client click on "Network Setup" in the context menu. It's open a really big dialog, but we are interest only the middle part like this:
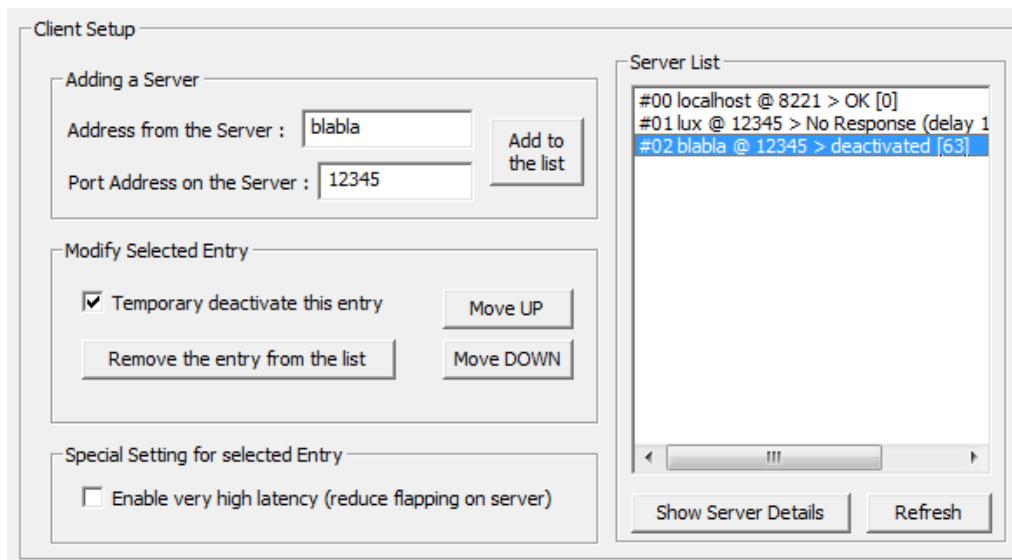
Adding a Server

Enter in the Address and Port field the values from the server. You can use IPv4, IPv6 or a hostname as address. By pressing the "Add to the List" button, your values insert in the next free server slot. The client will now try to ask the server about his data. After 2-3 seconds there should be some results. You can press the "Refresh" button to get a small status about the server. If the server not working select the server in the list and press the "Show Server Detail" to get more info what happen. In most cases the hostname can't resolve or the port number is wrong. You can't edit an entry, just delete it and enter a new one.
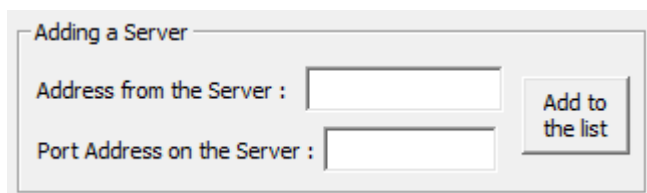
Please notice, the host name will only one time at start resolved. If you modify the DNS or hosts file you have to deactivate the server and activate it back again to renew the host resolve. This reduces the queries to name servers a lot. If you have a bunch of DHCP clients with changing IP, than you will

have problems. In this case contact me and I will implement another mechanism that fit your situation.

Modify Selected Entry

This section is activated and filled if you select a server in the list. The "Temporary deactivates the entry" checkbox immediate deactivate the server. The selected server will no longer ask for performance data. The "Remove the entry from the list" does exact what the button says. Removing a server includes removing the name resolution data and more stuff. This can take round about 1 second. The button changes the description while removing. And with "Move UP" or "Move DOWN" you can order the servers like you want. The ordering is the same order for the monitoring display. All changes you did take immediate effect, without asking, and store immediate into registry.
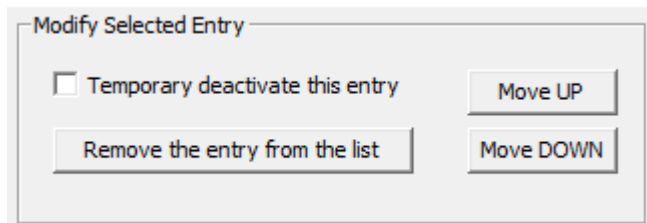
Special Settings for selected entry

This small section collects just one setting at this moment. With the "Enable high latency" checkbox you can enable a far server. The normal timeout for response is 100ms, that is enough for local connected hosts in local site LAN. Including some hops over routers. But if you try to add an internet server on the other side of the globe, you may have problems with flapping servers. They come and go or gone at the beginning. You can enable this setting to allow a very high latency of maximal one second reply. But this will immediate have impact to the other servers if the server really need that time or is simply offline. These can massive slowdown the collector. Your updating speed will drop down from regular one second down to unspecified. So take care about this setting.
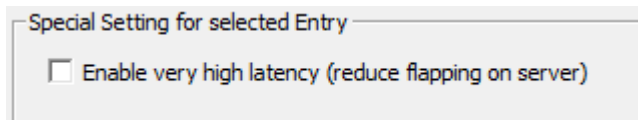
Server List

The server list contains the list of all servers including some important information. Please notice, this list is not automatically refreshing. Please take a look into next chapter to read why. The list entries start with a hash (#) follow by the entry slot number of the server. You can move entries up and down, see above. The next value is the hostname @ port you enter in the add server section. After the Greater-Than Symbol ">" comes the status of the server. Beware of brand new servers, they add sometimes with an unspecified status field, give him a second to add the server correctly and press refresh button. The values between the brackets "[xxx | yyy]" is the last query value and the current latency. See Version 8 – Current Latency Display for more info. On servers they are OK it should be 0, 1 or 2. This depends on the amount of servers and they summarized latency. The both buttons will be described in the next two chapters.
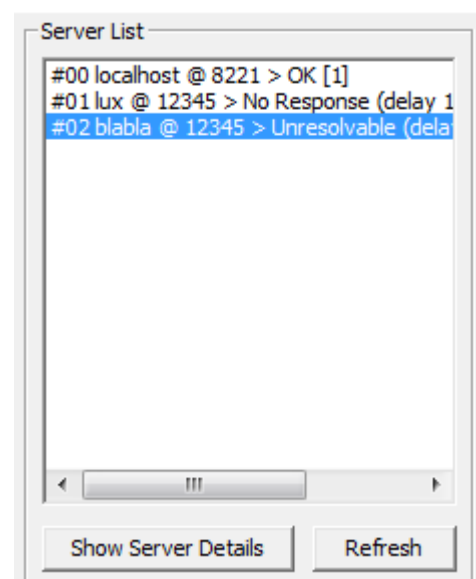
Refresh

The list is not automatically updated. Well I try to do this, but I fail. It was to strange to reselect and scrolling on the refresh, so I decide to drop 3000 lines source code and implement a static list with a refresh button. Sorry guys … but you have to live with that.

Show Server Detail

This function show a message box with all information about the server. The most information displayed already in the network setup dialog. The detail has additional information. One important is the status and last query. If you have many servers they didn't respond or have massive latency, it will slow down the collector process. Servers with status OK and last query of 3 or more seconds show you, you have a timing problem. Contact me and we can maybe do something.

Servers with the status UNRESOLVED query only every 120 seconds, and all other servers they are not OK queries every 15 seconds. Only servers with the status OK query each second.

Network Server Detail Dialog

Here some samples how they can look like (GUI overhead removed):
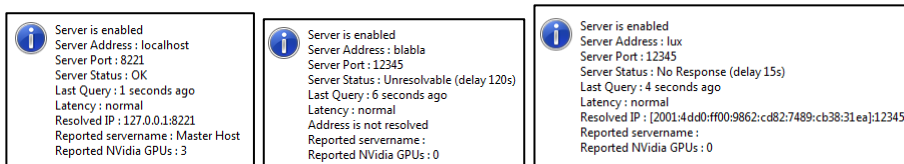
You can see all settings and info displayed well. The field "Last Query" shows you when the last activity of this entry happens. Unresolved status (second picture) where the hostname could not resolve, will delay 120 seconds. And on the 3th picture you see 15 second delay on no response servers. And you can see the resolved IP address that this server use. The server "lux" uses an IPv6 address. Anyway, if you have any problems, you can clearly find out what happen on this dialog.

Security

Introduction

Since this version can handle and answer requests from the network, there is always a change to make wired stuff. I really try to protect the processes and communication against remote exploits. There are several checks and safety code in there to prevent buffer overflows, buffer underruns, and any other things I didn't know. But this is not a guarantee that anything safe, nobody can give this.

If you use whNvInfo on public internet server, like render farm, you should use a firewall or packet filter to protect the running server from evil guys, if you don't trust me, but you can, I am a nice one.

I try to revoke any permission except these I really need from the process itself, but seriously the windows documentation about this need a full MIT exam to implement this. So I give up a little. If someone have some hints and sample codes, please give it to me, I could use it.

There are per default 3 threads running, one for the clients and two for the server (IPv4 and IPv6). They still active even no servers or clients used. This is a small helper hack to prevent an overkill of thread management and stuff. It makes it much more stable. If network server or no client, the thread never reach any network function, this means, if you turn off the network server, it is super safe against remote exploits.

If someone find a bug or can success remote exploit it, please give me a hint, I will fix it immediate.

To prevent a crazy much more bigger point of hack, whNvInfo did not have any auto update or version checks. This mechanism is the pure hell, because use of really unknown components while getting a resource across http or https. This is why it's not implemented. Best practice is you take a look every 5 minutes on my website to see if a new version online.

History

Version 12

Updating nvAPI to 410, updating Visual C++ Compiler to 2017. Hope try to fix the RPM Problem on RTX Cards. Updating Homepage URL.

Version 11

Bug fix the count of cores on the host. Reported by Lewis. The function I used counted only the cores of one processor group. This is now fixed. This has only effect on systems with more than 64 cores.

Version 10

Fixing a bug in NVAPI, the function NvAPI_GPU_GetAllClockFrequencies, returns sometimes a zero frequency on driver version 355.60. The fix detects a zero now and use simply the old value. That is not the best method, but it works for now. Furthermore I reduce the amount of functions calls to this function as well.

Version 9

Fixing a fan bug. It should now display the RPM, instead of percent.

Version 8

Modified client thread model

Each configured and activated server has now his own thread. This allows much higher latency on each server without blocking the rest. High Latency is now 10 seconds instead of 1 second, and the default latency is now 900ms. These thread running with low priority.

Current Latency Display

The current latency of the server is now display in the Detail message box of the server and in the list box. It is the value after the pipe symbol. The timer I use is a very simple timer and no high res timer. It is normal that he display zero, this happen if the server very fast. The resolution of the timer is between 10 up to 20 milliseconds. So don't worry about a zero.

Contact and Last Words

You can contact me per mail or twitter or in the starting thread where I announce whNvInfo by myself, like RedShift forum or some XSI forum. You also can use the comment function on my website too.

```
Wolfgang Hurst
whNvInfo@whurst.net
https://whurst.net
https://twitter.com/wolfgang_hurst
```

THANKS FOR READING, AND HAVE FUN WITH IT