



RuggedVPN Stable Firmware Release July 12, 2018 – Version 201805236/2018070900

This firmware release is a breakthrough: After working on it for three years, finally a brand new and highly improved implementation of our WAN optimization feature is back. It's much more compatible, faster and stable than its first incarnation from a decade ago.

Whatever link you might have, may it be bonded LTE, lossy DSL or even satellite: Our new firmware will perform better than ever. In addition to this, this firmware is bringing a metric ton of bug fixes and improvements, which you will find listed below. This release also contains a very important update if you are using Hub virtualization. Please note that this release does NOT fully support IPv6. Please check the release notes below for both issues.

An updated firmware image will be available on Amazon AWS as soon as their approval process is finished.

If you wish to upgrade from Classic firmware, please first update the router to the last stable Classic firmware release (Version 2015081830/2015102900 released on November 27th 2015). Please note that upgrading your firmware from Classic to RuggedVPN requires a Viprinet Lifetime Maintenance license to be active. For more information, please check <https://www.viprinet.com/vlm>. It is possible to have Routers and Hubs running on the latest version of Classic firmware connect to a device running RuggedVPN firmware. However, in this case a compatibility mode will be used, which limits performance and features. It is therefore not recommended to permanently use such a setup, but it is OK to have a Classic firmware device talk to a RuggedVPN firmware device while you are upgrading these devices. The Software VPN Client is available based on Classic Firmware and alternatively based on the RuggedVPN firmware generation. This is the final firmware version that still supports connecting old devices running our Classic firmware generation (2015 and prior) and upgrading from such a firmware release.

The list below lists all new features and bug fixes compared to the previous stable RuggedVPN firmware release (Version 2017102440/2018041200 released on March 24 2018).

New Features

- The WAN optimization feature is available as an option to be enabled in the QoS settings. If you restore your QoS Templates to manufacturing defaults, a new set of QoS classes including the WAN optimization feature will be available to be assigned to your tunnels.
- The download test tool has a new "delay" parameter that can be used to wait between sending the HTTP header and the actual data for a number of seconds – useful to simulate long-lasting idle connections. Like this: `wget -O NUL "http://192.168.200.1/exec?module=download&delay=10"`.
- The route back option in LAN settings now has an option to disable its logging to prevent excessive logging.

Bug Fixes

- Due to both virtualization hypervisors and the Linux kernel addressing a bug on how a virtual machine identity is passed on to the guest operating system, on multiple hypervisors the device might get into a state of "identity crisis" after a firmware update – it will identify as a new machine with a different hardware ID, not matching its virtual hub license. We hope to have fixed this problem with this release.

We've tested with latest versions of VMWare and KVM but if you are deploying Virtual Hubs please urgently check your own installations.

- Fragmented Packets are now also captured using the packet capture tool.
- Enabled Band 8 for 4.5G Europe/America.
- Fixed performance problems within IP fragmentation on 200/5xx.
- Fix for system getting stuck in a reboot loop if more than 256 routes are created.
- A big and complex fix for FEC:

The first problem was that FEC was meant as a "this is a minimum guaranteed redundancy for the system". This makes sense for parity, but for packet duplication it does not.

The logic was something like "If I want 4 channels, and in general have 4 channels connected / fitting the QoS needs, but only three of those can be used right now, then this means that we have reached the maximum available bandwidth and can not do anything".

Turns out that the speed tests after the missing channel reconnected caused exactly this scenario – all channels are there, but for this one channel all the potentially available traffic is eaten by speed tests.

The logic has now been changed: For parity it still is "we guarantee this minimum", while for duplication and up, now we just don't care – as long as we have a channel left, we'll send. This means that a potential QoS class doing lots of traffic could prevent a duplication class to only use a single channel most of the times. We still think it's more intuitive than it is today, and this limitation can be solved using bandwidth guarantees.

Also, we have fixed the situation that a speed test can take away all the traffic of a channel, not allowing much user traffic (and none in case of duplicate and above).

- "ICMP time exceeded" messages caused by a routing loop could itself cause a routing loop resulting in a packet storm.
- Fixed ICMP stuff so you again see the Viprinet hop when doing a traceroute.
- Lots of optimizations for 5xx throughput/CPU load.
- Fixed CPU Temperature Sensor for 50X0.
- Fixed internal error BEF3238723CD2983.
- Fixed a long-standing bug in RuggedVPN (has always been present) that in a rare case if multiple channels had the exact same score could have caused that only one of them would be used.
- The HTTP server will now correctly declare the character encoding (UTF8 / Ansi) for any kind of text file (html, js, css etc) served. This is in preparation for a fully Unicode-enabled web interface that can be localized in multiple languages. Please report any strange encoding/character errors.
- Fixed WLAN AP config regions now being selectable with Toughlink.
- Fixed the "HTTP Server is eating 100% of the CPU forever" problem.
- "No gratuitous ARP on stacking failover" bug should be fixed.
- Fixed the underlying problem leading to Internal Error 137239A239232341 / Map size HUGE. That one is actually interesting: If you did NOT use guaranteed delivery and also did NOT use WANOptimizer, over time any flow would eat a few bytes of memory until the flow ends. If you had VERY long running connections (say a VPN tunnel through the VPN tunnel), and/or connections that ran a long time with very small packets

(say a SIP trunk), the amount of memory available to the device could actually have mattered. Also, the longer this flow lasted, (a little) more CPU would have been used.

- All Gratuitous ARPs will be repeated after 5 seconds in case the first one gets lost for whatever reason. This fixes a problem with ARP on stacked nodes.

Known Issues

- We are not happy with the maximum bonding throughput we are seeing on the 2610 and 2620. This will significantly improve with the next firmware release. If you are interested in trying a beta, let us know.
- IPv6 support had a lot of really bad bugs, one of which could cause very high CPU load. To be frank, IPv6 support has been broken since last December. As nobody complained after the last two stable releases, we assume it's not a big problem to our customers. Therefore, for this stable release we are now dropping most problematic IPv6 traffic. After this stable release, we'll re-enable IPv6 support after having fixed the bugs. If anyone of you is interested in testing IPv6 setups, let us know.
- If you have both WANoptimizer and non-wan-opt connections in a QoS class (say WANopt is enabled within QoS, but the same class is used for UDP), there might be starvation of non-wan-opt traffic. Please try not to have both UDP and TCP traffic in a QoS class where WAN optimization is enabled.
- There might be problems with fixed-bandwidth applications (Video streams) using the WANoptimizer. Please report back if you see any.
- The 300 is very low on memory. If you want to update/downgrade the firmware after using this release, you need to first reboot the device to have enough memory. In general, WANoptimizer is of limited use on the 300 due to the low amount of RAM it has. We highly recommend using one of our 300-to-310 trade-in offers and get rid of this outdated device.
- On 200 and 5xx, if you try to reach Integrated Services (web interface, ping to LAN IP, CLI etc) you will see packet loss and/or high ping times if the router is idle. If the router has some load, the problems disappear.

This is a side-effect of optimizing WANoptimizer performance for these products and can not easily be fixed. Therefore it will stay this way for this release, and it will be fixed during the next beta cycle.

In real life, this should not be much of an issue – however, it may confuse your monitoring systems (in our case, it did confuse an automated test doing a CLI login on an idle router).