

STATELESS TRANSPORT TUNNELING (STT)

STT is another tunneling protocol along the lines of the VXLAN and NVGRE proposals. As with both of those the intent of STT is to provide a network overlay, or virtual network running on top of a physical network. STT was proposed by Nicira and is therefore not surprisingly written from a software centric view rather than other proposals written from a network centric view. The main advantage of the STT proposal is it's ability to be implemented in a software switch while still benefitting from NIC hardware acceleration. The other advantage of STT is its use of a 64 bit network ID rather than the 32 bit IDs used by NVGRE and VXLAN.

The hardware offload STT grants relieves the server CPU of a significant workload in high bandwidth systems (10G+.) This separates it from it's peers that use an IP encapsulation in the soft switch which negate the NIC's LSO and LRO functions. The way STT goes about this is by having the software switch insert header information into the packet to make it look like a TCP packet, as well as the required network

virtualization features. This allows the guest OS to send frames up to 64k to the hypervisor which are encapsulated and sent to the NIC for segmentation. While this does allow for the HW offload to be utilized it causes several network issues due to its use of valid TCP headers it causes issues for many network appliances or “middle boxes.”

STT is not expected to be ratified and is considered by some to have been proposed for informational purposes, rather than with the end goal of a ratified standard. With its misuse of a valid TCP header it would be hard pressed for ratification. STT does bring up the interesting issue of hardware offload. The IP tunneling protocols mentioned above create extra overhead on host CPUs due to their inability to benefit from NIC acceleration techniques. VXLAN and NVGRE are intended to be implemented in hardware to solve this problem. Both VXLAN and NVGRE use a 32 bit network ID because they are intended to be implemented in hardware, this space provides for 16 million tenants. Hardware implementation is coming quickly in the case of VXLAN with vendors announcing VXLAN capable switches and NICs.

Source: <http://www.definethecloud.net/stateless-transport-tunneling-stt/>