

SFD17 - StarWind

\*Vlad Karaiev

3 main directions

- SDS - StarWind Virtual SAN
- HCI and Storage
- VTL

Strict focus on SMB and Enterprise ROBO

Milestones

2002 - iSCSI to Windows

2008 - shift to HCI

2014 - shift to turnkey HCI

2016 - iSER for Windows

2018 - NVMoF to Windows and vSphere

\*Max

NVMe over Fabrics - the next big thing in [network] storage

[StarWind\_MaxPresents.png]

SAS - designed for disk

NVMe - designed for Flash

[StarWind\_SASvsNVMe-oF.png]

Native Target for Windows v1

Windows Server (no HyperV)

High performance I/O

- use CPU for everything
- No Kernel <-> User jumps
- No context switches

CPU is a cheap resource

Problem - no SPDK in Windows, reduced kernel not no kernel

Result - overhead is still high

\*Demo 1 with Taras

Native target for Windows v1 (with native Linux initiator)

Without SPDK one NVMe drive performs OK, but it eats 8 CPU cores alive. If we want to scale we need 100% polling in user mode.

Now with v2 - software moved into a VM (HyperV with Linux NVMf initiator)

100% user mode

\*Demo 2 with Taras

(Sneak preview of StarWind Stack mgmt interface)

“We found the elegant way to bring missing SPDK functionality to Windows Server:  
Run it in a VM with proper OS!  
First benefit - CPU is used more efficiently”

\*Demo 3 with Taras  
Also on vSphere

\*Max  
NVMe-oF initiator for Windows  
Combined Linux Initiator and stub driver in hypervisor to get properly working initiator  
Benefits  
- performance  
- portability (90% shared code)  
- stability through isolation (jails)

Next step is to do the same thing with vSphere

\*Vlad - Use Cases  
Data Centres (including big ones like AWS and Azure)  
Goal: deliver lossless flash storage to the VM or application that needs it

SMB  
Goal: Aggregate flash and improve ROI in brownfield environments  
Work with Microsoft S2D or VMware vSAN  
MS Initiator available mid-October, VMware support planned for December

StarWind HCA and SA  
Goal: East-West storage traffic (currently using iSER)

NVMe-oF  
What about HA and multipathing?  
Want to deliver lossless to initiator (in terms of performance)

Software or appliances

Anton “CPU is the cheapest resource we have right now”