



Windows Server® 2008 R2

Linux on Hyper-V– Our Journey Through The Linux Staging Tree

*Dr. K. Y. Srinivasan
Principal Architect
Microsoft Corp
Redmond, WA*

*Tom Hanrahan
Principal GPM
Microsoft Corp
Redmond, WA*



Windows Server® 2008 R2

By The Numbers

Initial Release

- Linux Hyper-V drivers released to the community on July 29, 2009 under GPLv2 License
- Initial release consisted of four drivers:
 - vmbus, storvsc, blkvsc, netvsc,
- Two drivers soon followed
 - utils, timesource
- Initial location in the kernel tree: `drivers/staging/hv/`
- Total number of lines of code: About 21,000



Serena Service Manager



All About Microsoft

Mary Jo Foley

- Mobile
- RSS
- Email Alerts

Home / News & Blogs / All About Microsoft

Pigs do fly: Microsoft unleashes 20,000 lines of Linux code

By Mary Jo Foley | July 20, 2009, 8:58am PDT

Summary: *Nope, that's not a typo in the headline. Microsoft is releasing three Microsoft-developed Linux drivers to the Linux community for possible inclusion in the Linux source tree, the company announced on the opening day of the OSCON open-source conference.*


211
Comments

2

Tweet

0

The Effort To Exit Staging (Patches)



Community Participation

- In addition to Greg, following community developers contributed to our drivers:
 - Stephen Hemminger (26 patches)
 - Bill Pemberton (68 patches)
- Other community help: reviews etc.
 - Dan Carpenter
 - Joe Perches
 - Sasha Levin
 - Olaf Herring
 - Christoph Hellwig

Key Milestones

- Linux 3.0:
 - With 361 Patches MSFT was the 7th highest contributor
 - With 343 Patches a MSFT engineer was the top contributor
- Linux 3.2:
 - With 177 Patches MSFT was the 19th highest contributor
 - With 165 Patches a MSFT engineer was the 5th highest contributor

Key Dates

- Oct. 4, 2011
 - Timesource driver exits staging
- Oct. 11, 2011
 - vmbus and utils drivers exit staging
- Nov. 26, 2011
 - Mouse driver exits staging
- Nov. 29, 2011
 - Netvsc driver exits staging
- Feb. 13, 2012
 - Storage drivers exit staging

Top Five Linux Contributor: Microsoft

By Steven J. Vaughan-Nichols | July 17, 2011, 1:45pm PDT

Summary: *While you shouldn't expect Windows to be open-sourced in your life-time, Microsoft—yes, Microsoft—is the fifth largest code contributor to Linux 3.0.*

Cats and dogs; apples and oranges; Linux and Microsoft. Two of these three things do not go together. Would you believe that Microsoft—yes Microsoft—was the fifth largest contributor to the soon to be released [Linux 3.0 kernel](#)? Believe it.

In a [Linux Weekly News](#) story, currently only available to subscribers, an [analysis of Linux 3.0 contributors](#) reveals that Microsoft was the fifth largest corporate contributor to Linux 3.0. While only 15h overall, that still puts Microsoft behind only Red Hat, Intel, Novell, and IBM in contributing new code to this version of Linux.



To be exact, Microsoft developer K. Y Srinivasan gets the credit for helping to improve Linux. Of course, as you might guess, neither Srinivasan nor Microsoft are doing this due to any particular love for Linux per se.

The vast bulk of Microsoft's contributions has been to its own [Hyper-V](#) virtualization hypervisor drivers. Hyper-V is Microsoft's 64-bit hypervisor-based virtualization system. It's Microsoft's answer to VMware and Linux's own native [Kernel-based Virtualization Manager \(KVM\)](#).

Passage Through The Staging Tree

- The process has benefitted our drivers immensely:
 - The drivers conform to the Linux Driver Model
 - The code is significantly cleaner and tighter
 - Excess of 60% reduction in the overall lines of code
 - Significant performance and stability improvements in all the drivers
 - Solid architectural base from which to enhance our support of Linux on our platform

Lessons Learned

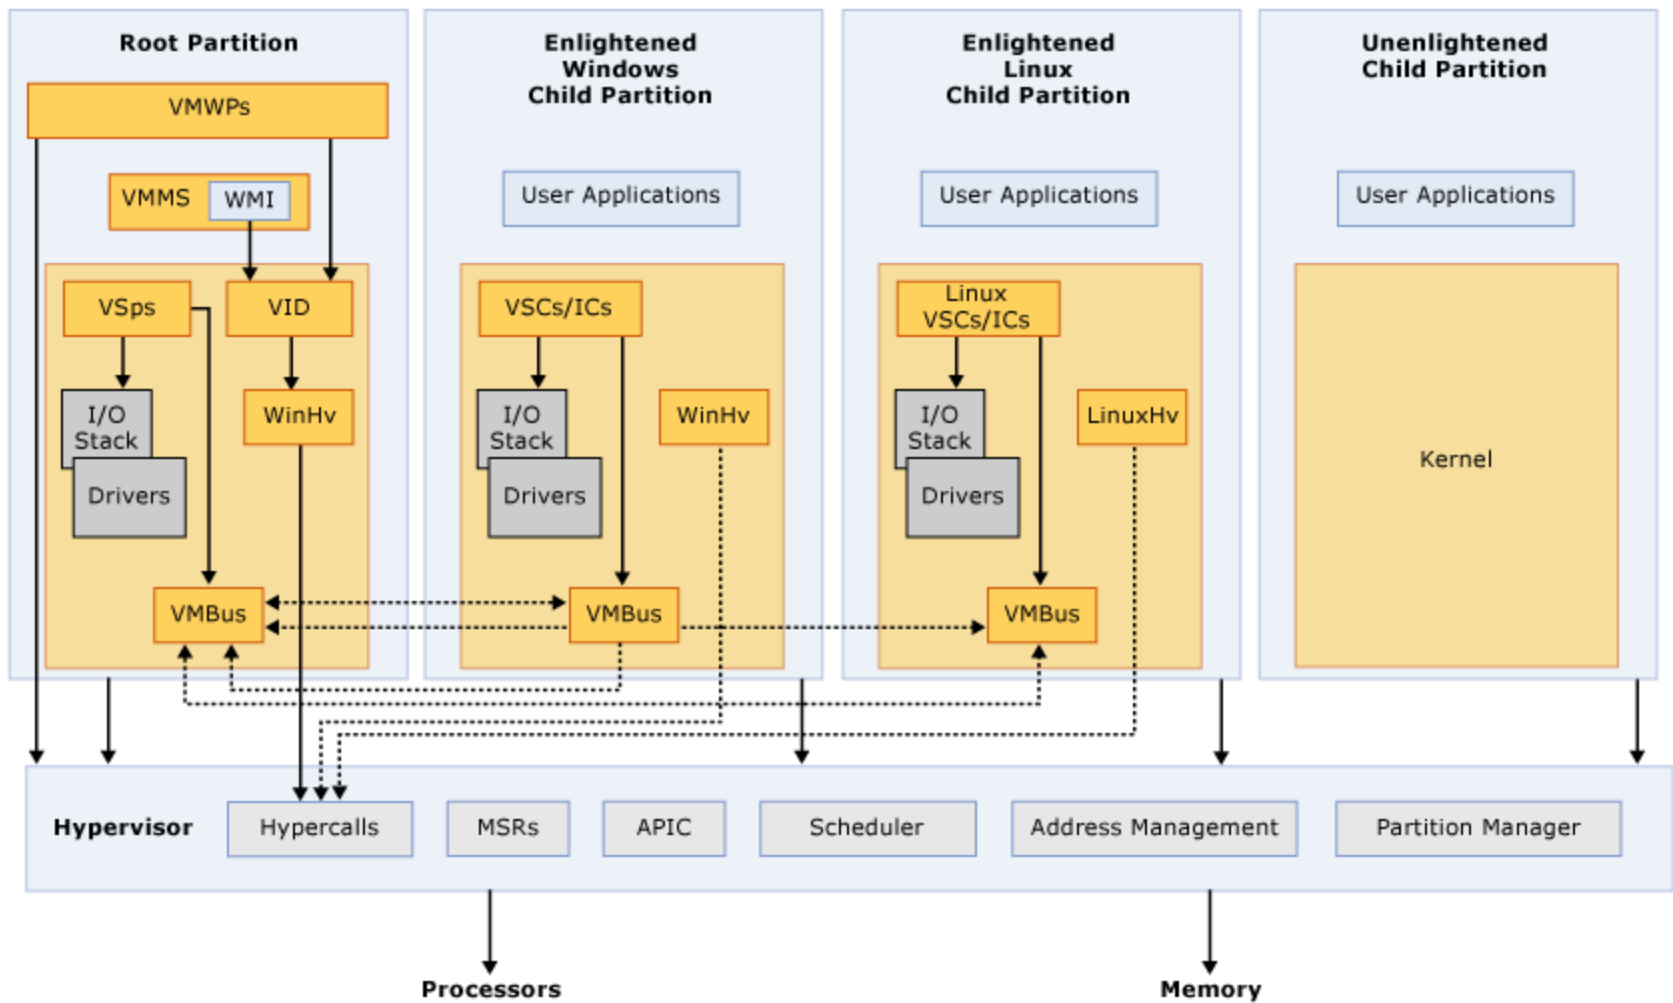
- Staging tree is not designed for quick (massive) re-architecture/cleanup efforts
- Upstream developers do not care for company specific product deadlines
- Need to be extremely responsive with the community



Windows Server® 2008 R2

Driver Updates

Hyper-V High Level Architecture



Hyper-V Architecture

- Hyper-V is a Type-1 Hypervisor for the x86 Instruction Set Architecture
- Requires Hardware Assist for Virtualization
- Emulates a standard x86 platform for guest Operating Systems
- Supports both 32-bit and 64-bit Guest Operating Systems

Hyper-V Architecture

- Full Virtualization with selective enlightenments:
 - Enlightened I/O Paths
 - Other low-level enlightenments
 - Time keeping
 - Context switching
 - TLB shoot-down etc.

Linux On Hyper-V

- Currently Linux hosted as a Fully virtualized guest with I/O enlightenments:
 - Standard kernel binaries supported
 - I/O enlightenments packaged as driver modules
- Linux pvops framework can be used to leverage additional Hyper-V specific enlightenments

VMbus Driver – hv_vmbus

- This is a Bus driver as far as the guest is concerned:
 - An ACPI Bus driver
 - Manages communication with the host on a per channel basis
 - Manages the life-cycle of channels
 - All “enlightened” guest devices belong to this virtual bus

The Hyper-V Channel


- GUID-based naming
 - Instance GUID
 - Type GUID
 - An open channel represents an “enlightened” device in the guest
- Each Channel represents a “service offer” from the host
- A pair of ring-buffers for communication with the host

Storage Driver – hv_storvsc

- Handles both SCSI GUID as well as IDE GUID
 - All block devices are handled as SCSI devices
- Implements a virtual HBA “front-end” driver
 - Registers itself as an HBA with the guest OS
- Currently does not support devices other than disk devices
- Supports hot add/remove SCSI devices

Storage Driver – hv_storvsc

Bonnie++ Block Write




Network Driver – hv_netvsc

- Implements a Synthetic NIC device for the guest:
 - Structured as a remote NDIS driver: NDIS protocol is used to communicate with the host
 - Implements the mapping between NDIS packets and SKBs
 - Copy free on the send path
 - Currently one copy on the receive side
 - Zero-copy implementation underway

Network Driver – hv_netvsc

Network Performance via Iperf



Utilities Driver – hv_utils

- The utils driver supports the following services:
 - Graceful Shutdown
 - Timesync
 - Heartbeat
 - Key/Value Pair (KVP)
 - Fully qualified domain name
 - IP address
 - OS build information
 - Linux Integration Services version

Clock Source Driver – hv_timesource

- Provides a pluggable time-source for the Linux guest:
 - Based on Hyper-V supported per-partition counter
 - Provides a stable clock source for the Linux guest

Mouse Driver – hv_mousevsc

- This is a “hid compliant” mouse driver for Linux guests on Hyper-V:
 - Supports seamless mouse tracking on the server console
 - Supports seamless mouse tracking on a Remote Desktop session

Linux Guests Supported

- Fully integrated with SLES10, SLES11 and Ubuntu 12.04
 - Hyper-V drivers are part of the installation media
- Drivers are available to support RHEL/CentOS 5 and RHEL / CentOS 6
 - RHEL 6 supports all features
 - RHEL 5 does not support mouse or KVP
- Community members have deployed our drivers on other distros as well

