


BROCADE



FCoE & Server Provided MAC Addresses

Single OS Systems and Virtualizing Systems

T11/07-591v1

John L. Hufferd

The basic assumption of these slides

The important position that is being taken in these slides is:

→The FCoE adapter should be seen at the Host system as a normal FC adapter

- o Investment protection
- o Evolutionary convergence

We want the HBA interface to be seen as “Vanilla” FC

(There could, over-time, be some vendor special management functions, which each vendor will claim is better than every other vendor but the Driver will start out as, and continue to be, a FC Driver.)

If all FC functions can be accomplished via one Real FC N_Port (NPIV, FLOGI, FDISC, etc.), then an FCoE adapter which acts as a FC adapter and uses one “Burnt-in” MAC address can accomplish the same things



Basic position regarding MAC addresses

→ No changes to the way servers define & use MAC address

- o Can come from “Burnt-in” MAC Addresses
- o Can come from Server Created Virtual MAC Addresses
- o May be administratively configured/allocated by management modules

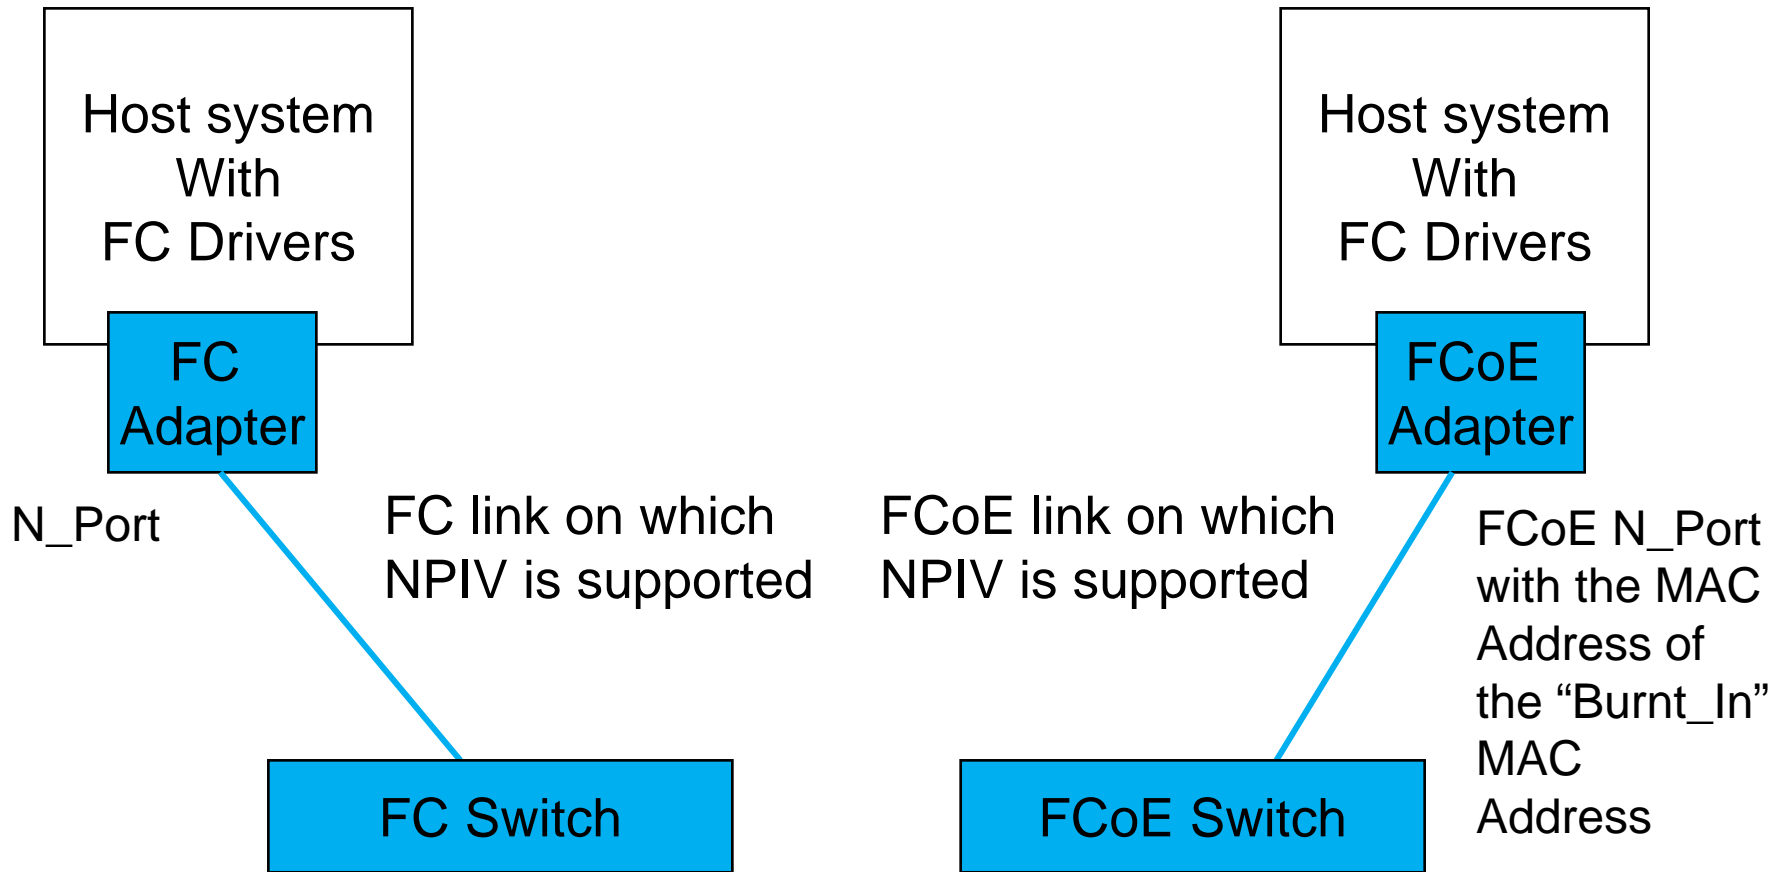
→ FCoE should not change the way Servers (and customers) manage MAC addresses today

- o Investment protection
- o Evolutionary convergence

→ Server Provided MAC Addresses ←



One to one comparison of Real FC and an FCoE link using a single MAC Address



Note about Discovery

These slides do not go into “discovery”

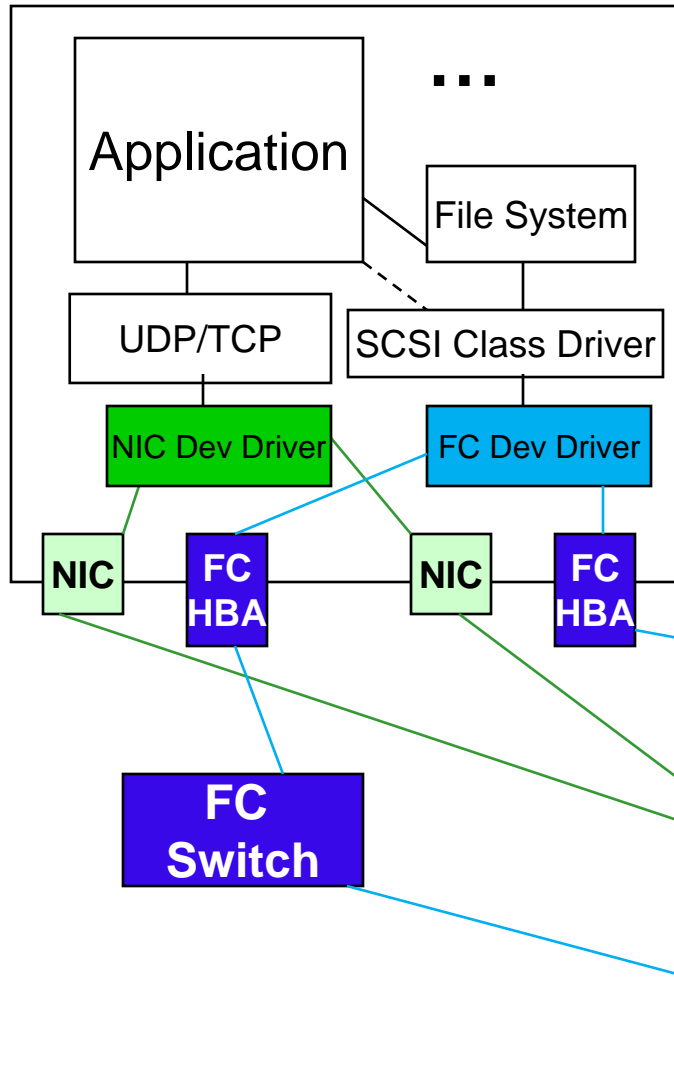
Discovery is a separate topic that needs its own set of discussions



Single OS systems

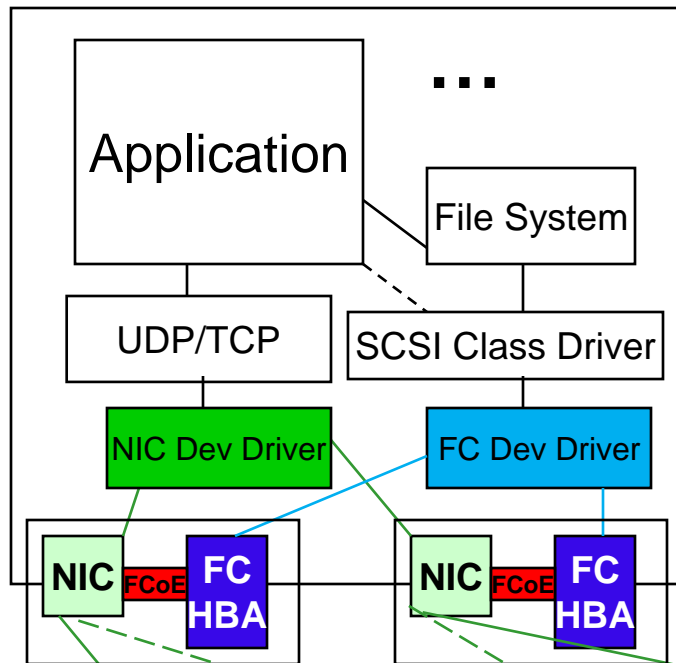


Single OS systems (Ethernet and FC)




- The Network support and Drivers will normally support a NIC's "Burnt-In" Universal MAC address
- Port Binding and other specialized functions may use a Locally created Address, or one of the additional "Burnt-In" MAC Addresses
- Many/Most NICs come with the capability for 4 or more "Burnt-In" MAC addresses

Single OS systems (Ethernet and FCoE)




- FCoE adapters, will probably still have host interfaces for both FC and NIC sides
- The FC function combined with the FCoE encapsulation function can use the same NIC as the rest of the system
- The “Burnt-In” MAC address should be just fine for this action since a single MAC address is the same concept as a single FC port
 - Discrimination based on Ethertype
- Many/Most NICs come with the capability for 4 or more “Burnt-In” MAC addresses, and if needed the implementation can chose any one of these (but this also works if only one MAC Address is implemented)

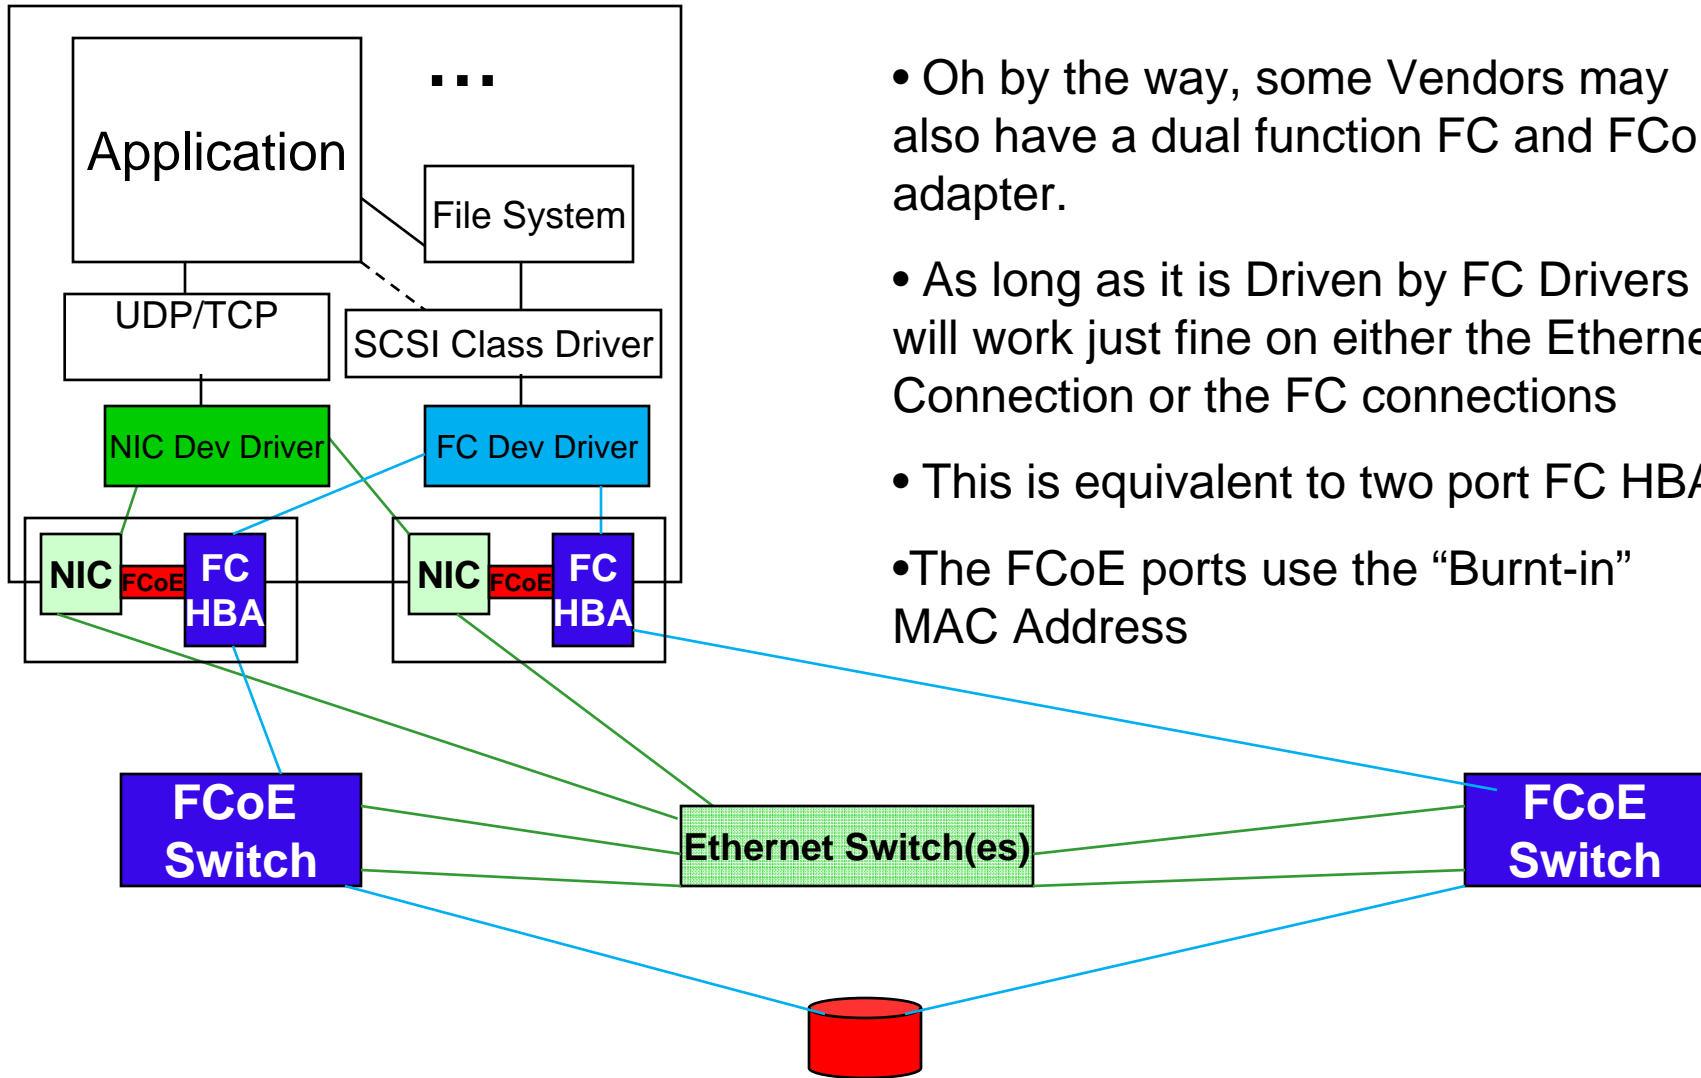


FCoE Switch can also include normal Ethernet Switch Function

So a separate Ethernet switch may not be needed



Single OS systems (Dual Function Ethernet/FCoE and FC)

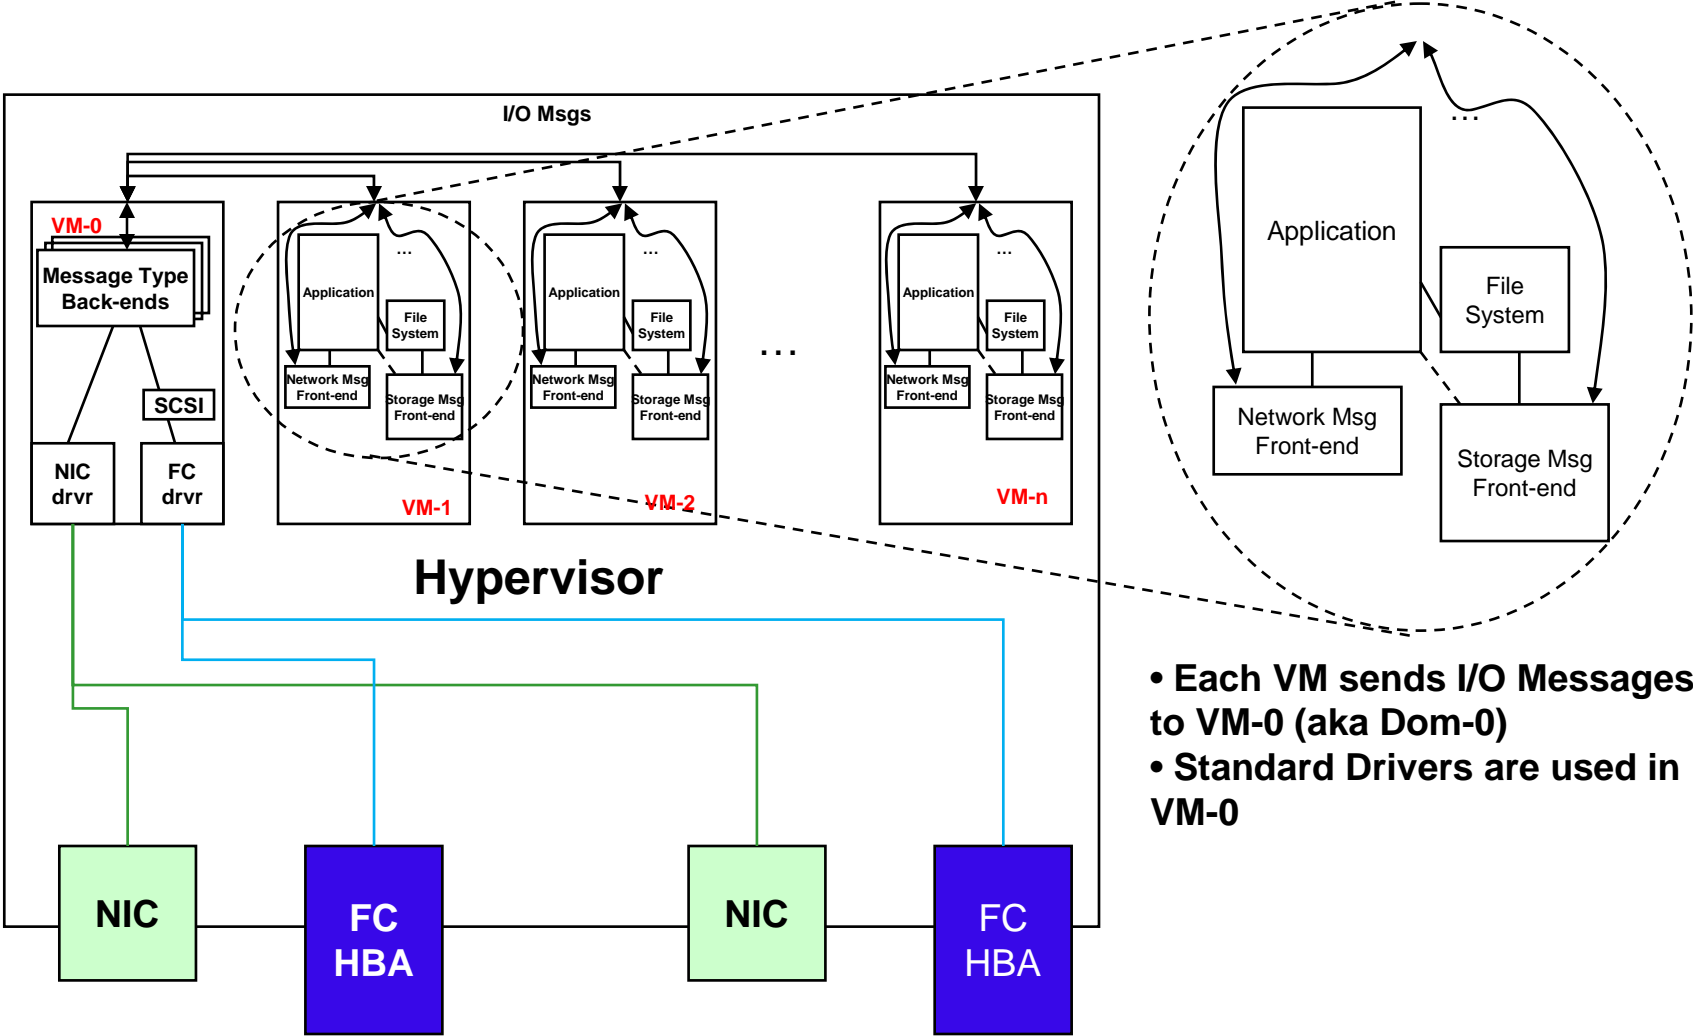


- Oh by the way, some Vendors may also have a dual function FC and FCoE adapter.
- As long as it is Driven by FC Drivers it will work just fine on either the Ethernet Connection or the FC connections
- This is equivalent to two port FC HBAs
- The FCoE ports use the “Burnt-in” MAC Address

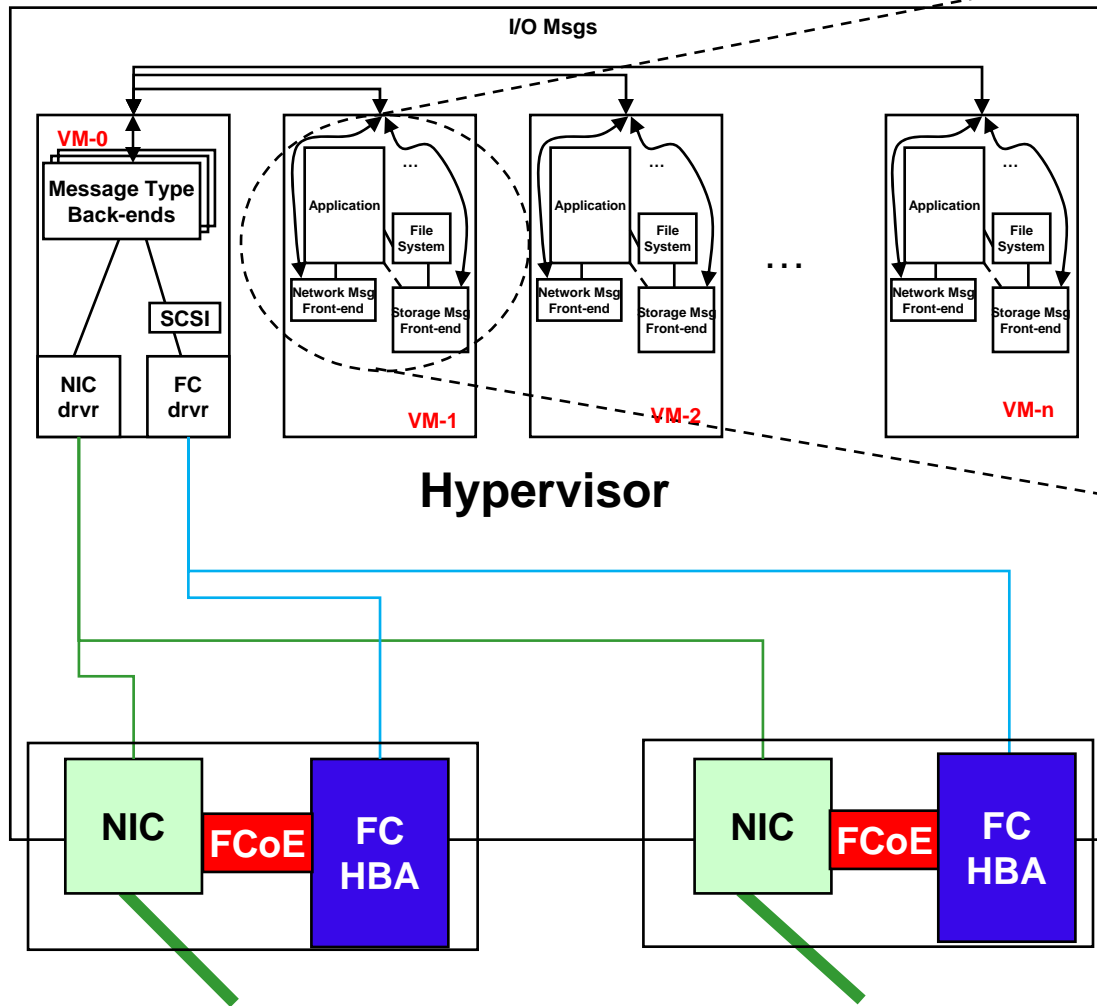
Virtualizing Environments



Para Virtualization (as seen in Xen)



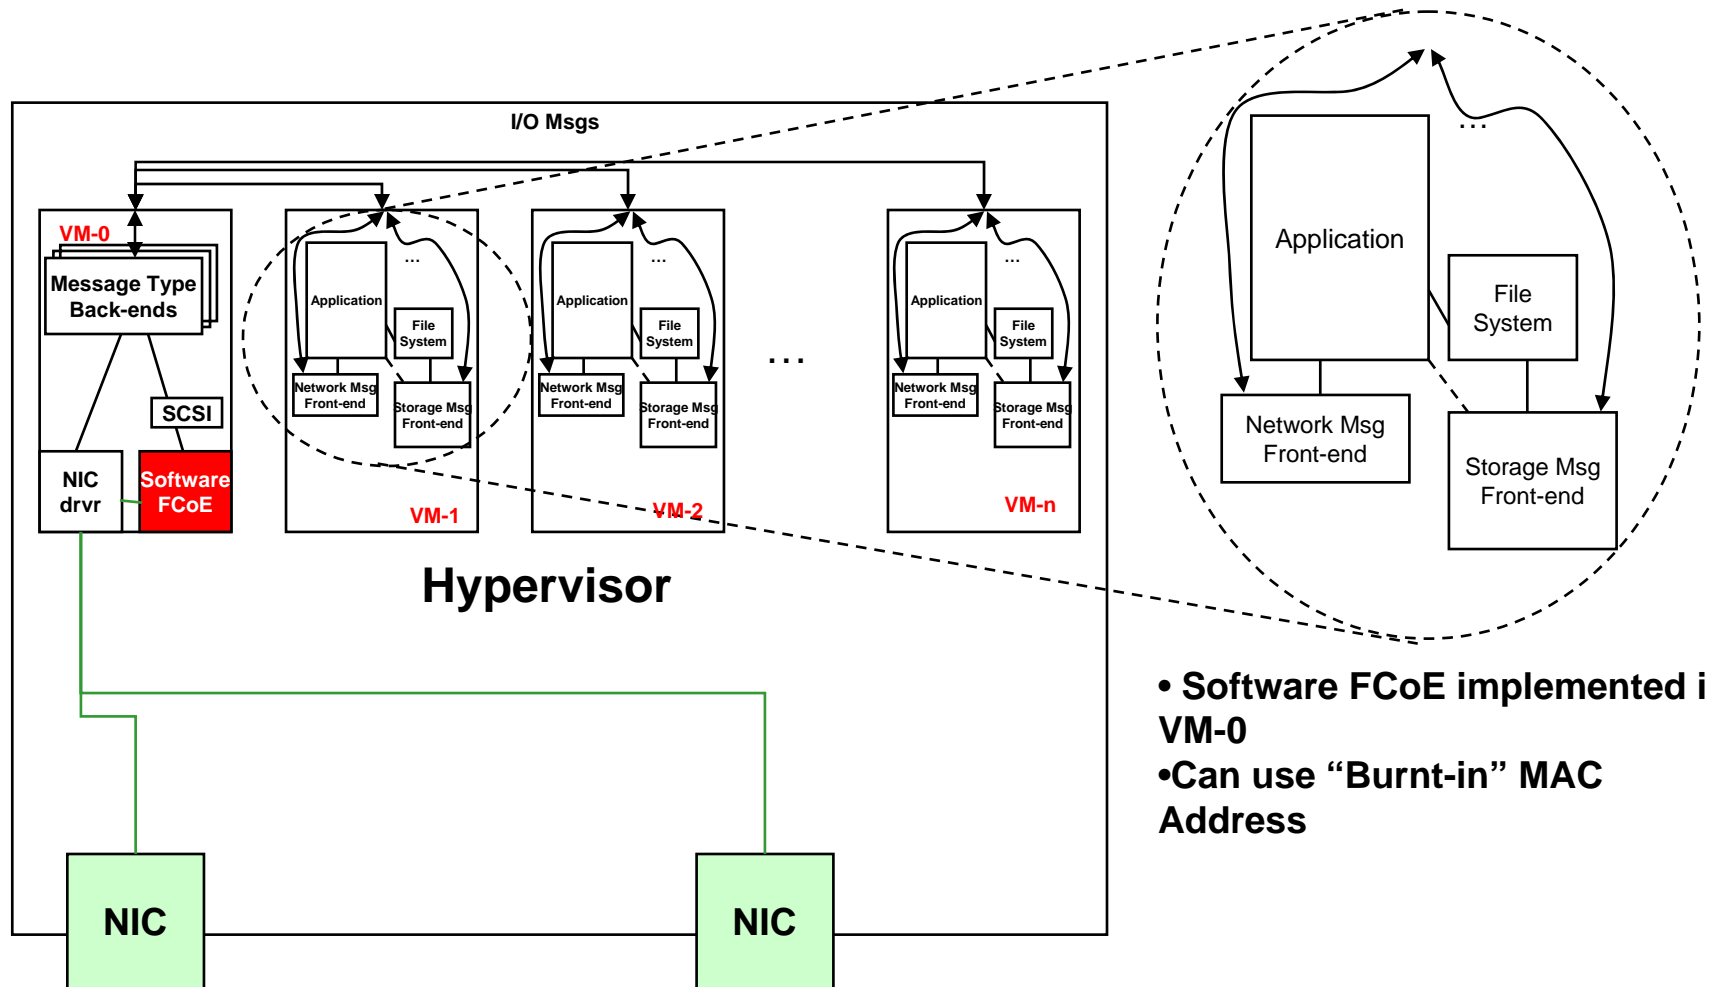
Para Virtualization (with FCoE)



- Continues to use the real FC Device Drivers
- Does not get involved in the use of the “Burnt-in” MAC Address used by the Adapter
- The same real FC NPIV capable drivers are used for FCoE

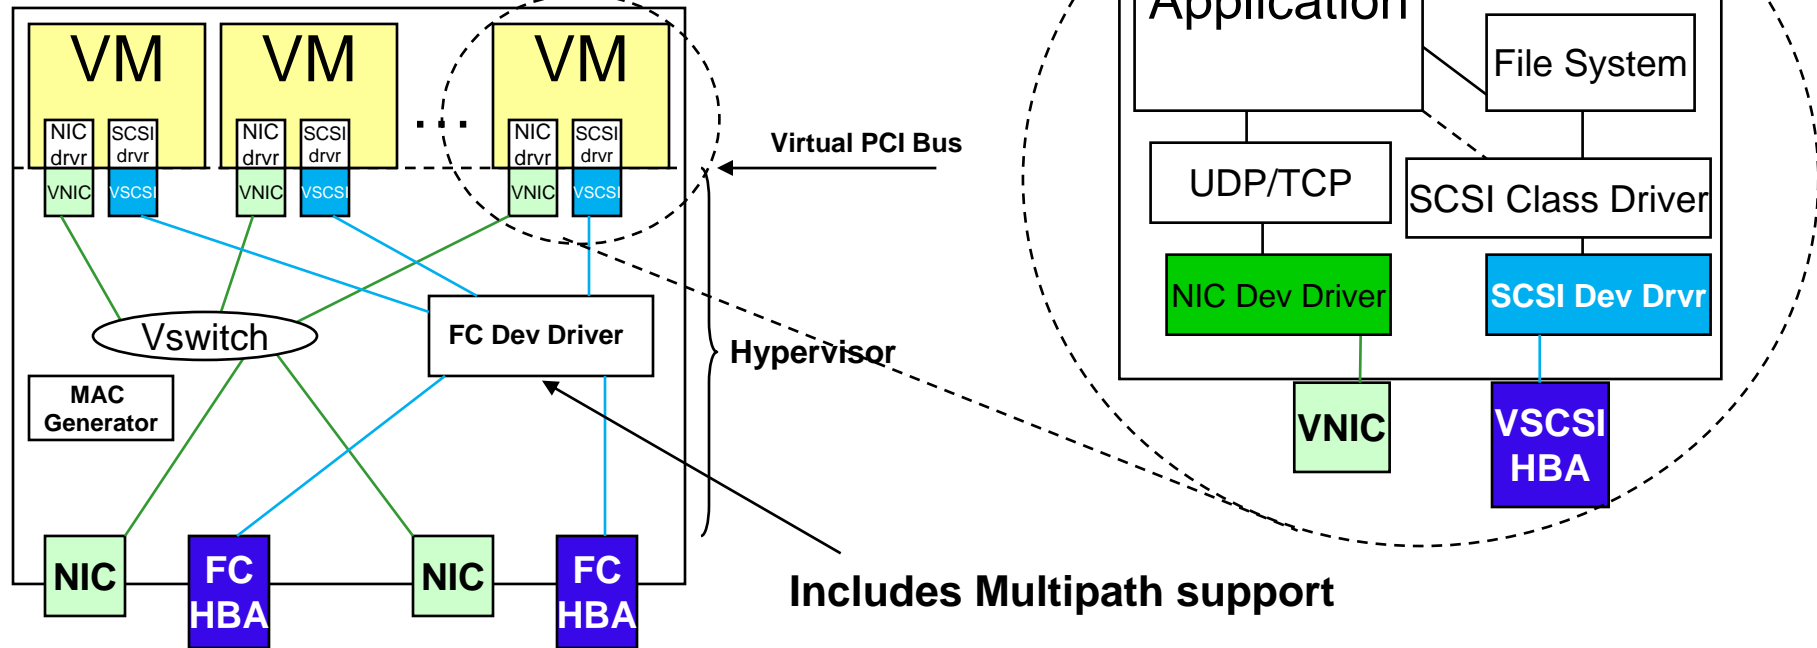


Para Virtualization (& Software FCoE)



Hardware Virtualization (as seen in VMware)

Each VM sees a Virtual Motherboard
With Virtual NICs and Virtual SCSI HBAs

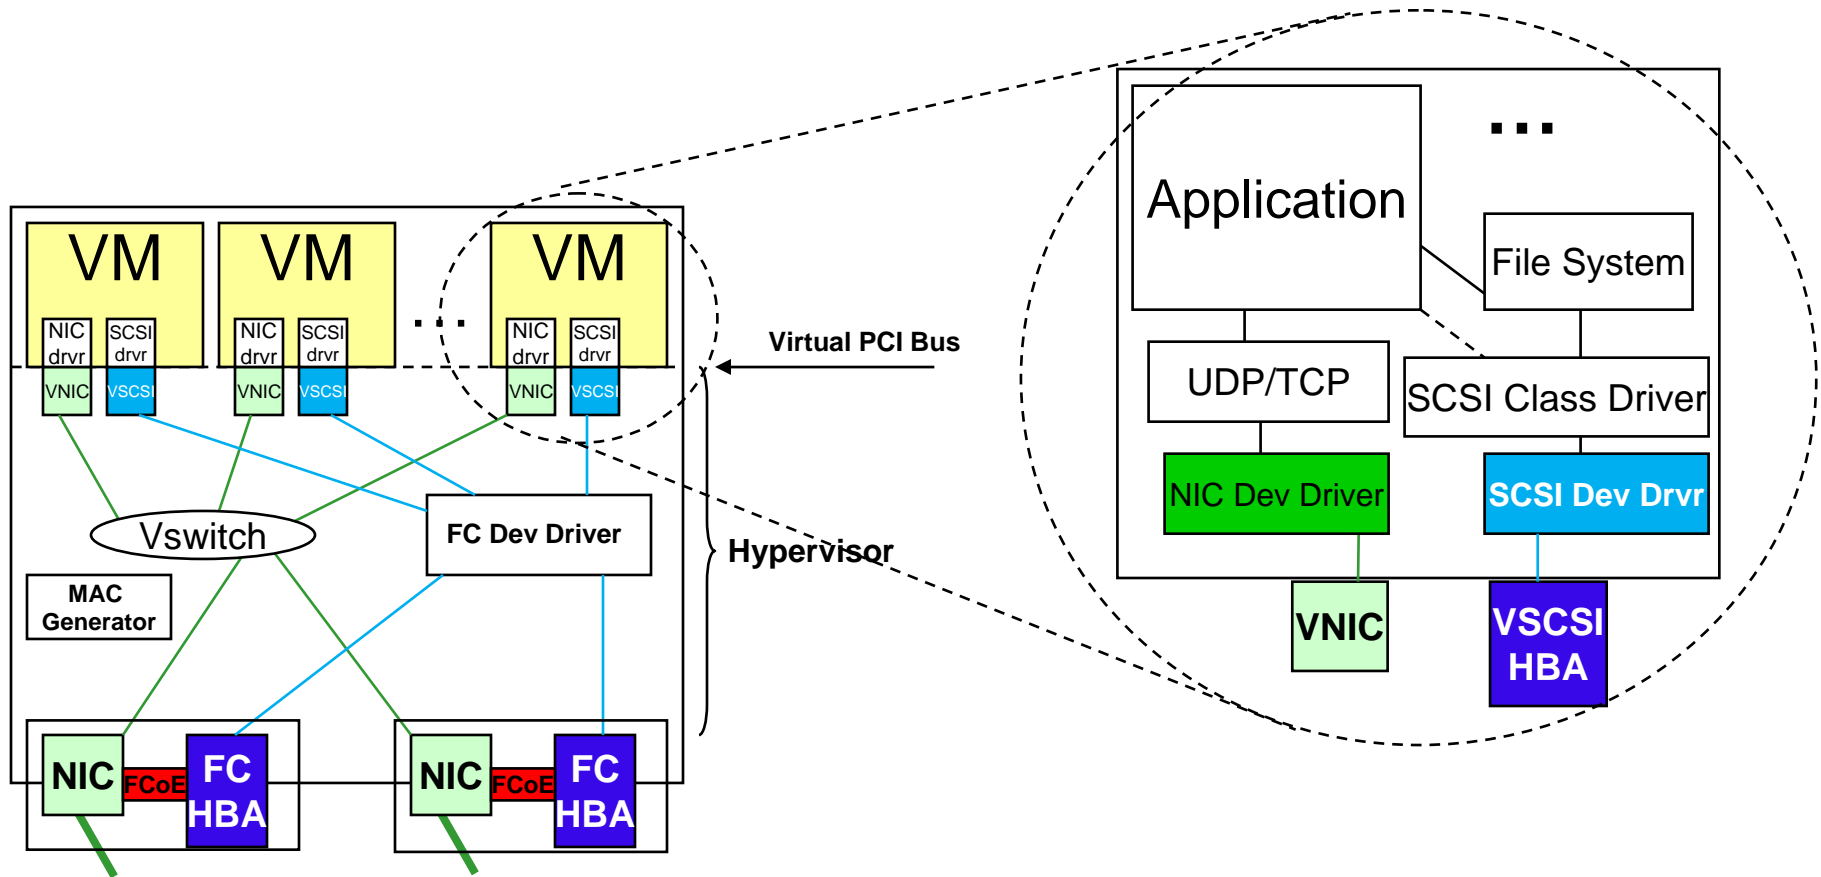


• **The Hypervisor:**

- Contains FC Device Driver & Virtual Switch to connect the VNICs and Real NICs
- Contains a MAC Address Generator that assigns MAC Addresses to each VNIC
 - MAC ADDR = OUI1 || (HASH (UUID) + Increment as needed) or OUI2 || (Admin setting)
 - (VMware UUID is based on the path to the virtual machine and the host's SMBIOS UUID)
- Uses FC HBA WWN for non NPIV support & Generates WWNs for NPIV support



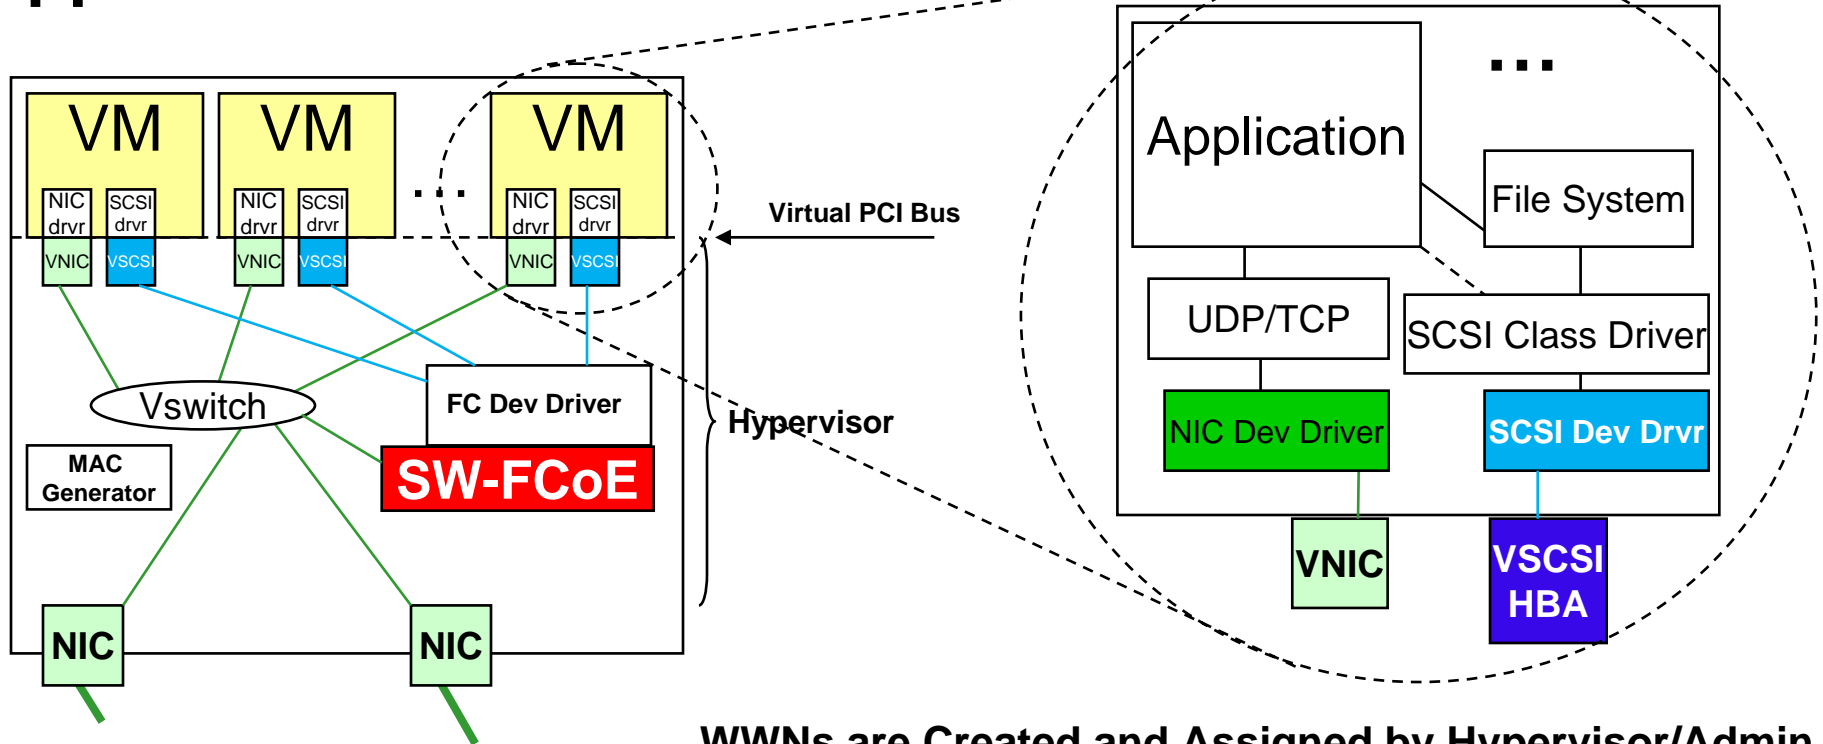
Hardware Virtualization (with FCoE HBAs)



- **The Hypervisor:**
 - **Continues to use the real FC Device Drivers**
 - **Does not get involved in the use of the “Burnt-in” MAC Address used by the Adapter**
- **The same real FC NPIV capable drivers are used for FCoE**



Hardware Virtualization (with possible Software FCoE) Approach 1



- **The Hypervisor:**

- Can use a FC Device Driver with additional “software FC HBA” with FCoE
- In an example of Server Provided MAC Addressing

- ✓ May use the “Burnt-in” MAC Address

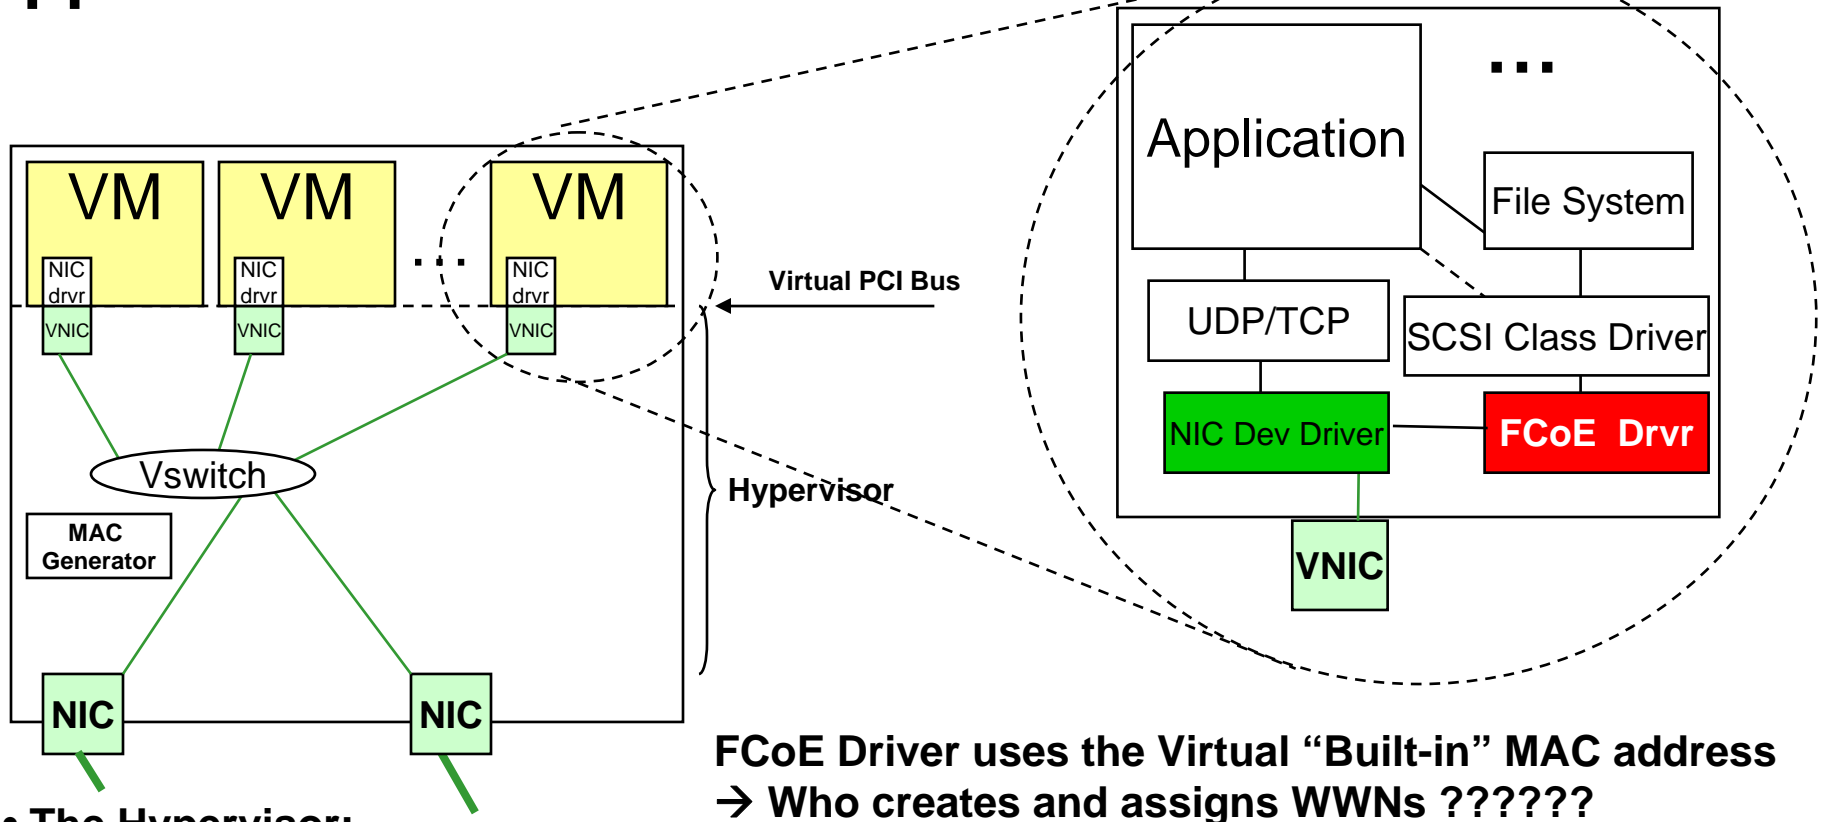
Or

- ✓ May use the MAC Address supplied by the Hypervisor MAC Address Generator or Admin

- **FC NPIV capable drivers can be used for FCoE**

WWNs are Created and Assigned by Hypervisor/Admin

Hardware Virtualization (with Possible Software FCoE) Approach 2



- **The Hypervisor:**

- **Treats Storage I/O just like Network I/O**

- (may be a problem with Vmotion – see later slides)

- **Assigns the Virtual NIC MAC Address as usual (Server Provided MAC Addressing)**

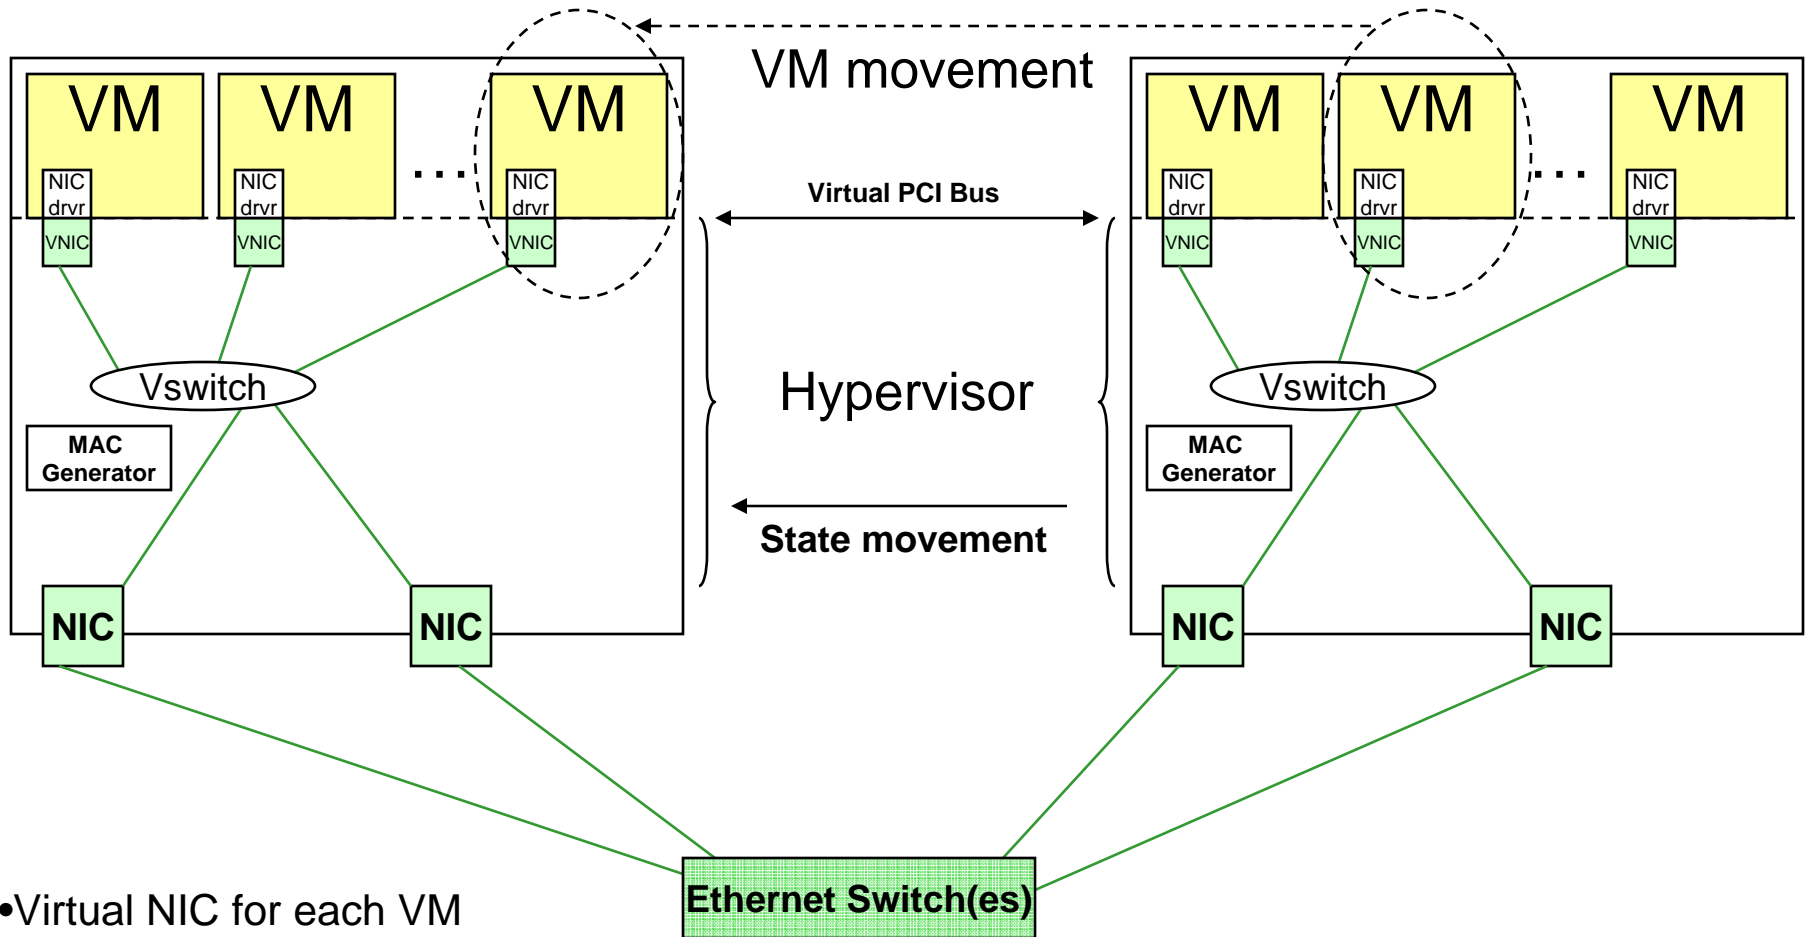
- (may have a separate VNIC for Storage)

- **NPIV support not needed**

VMware and Vmotion Environment

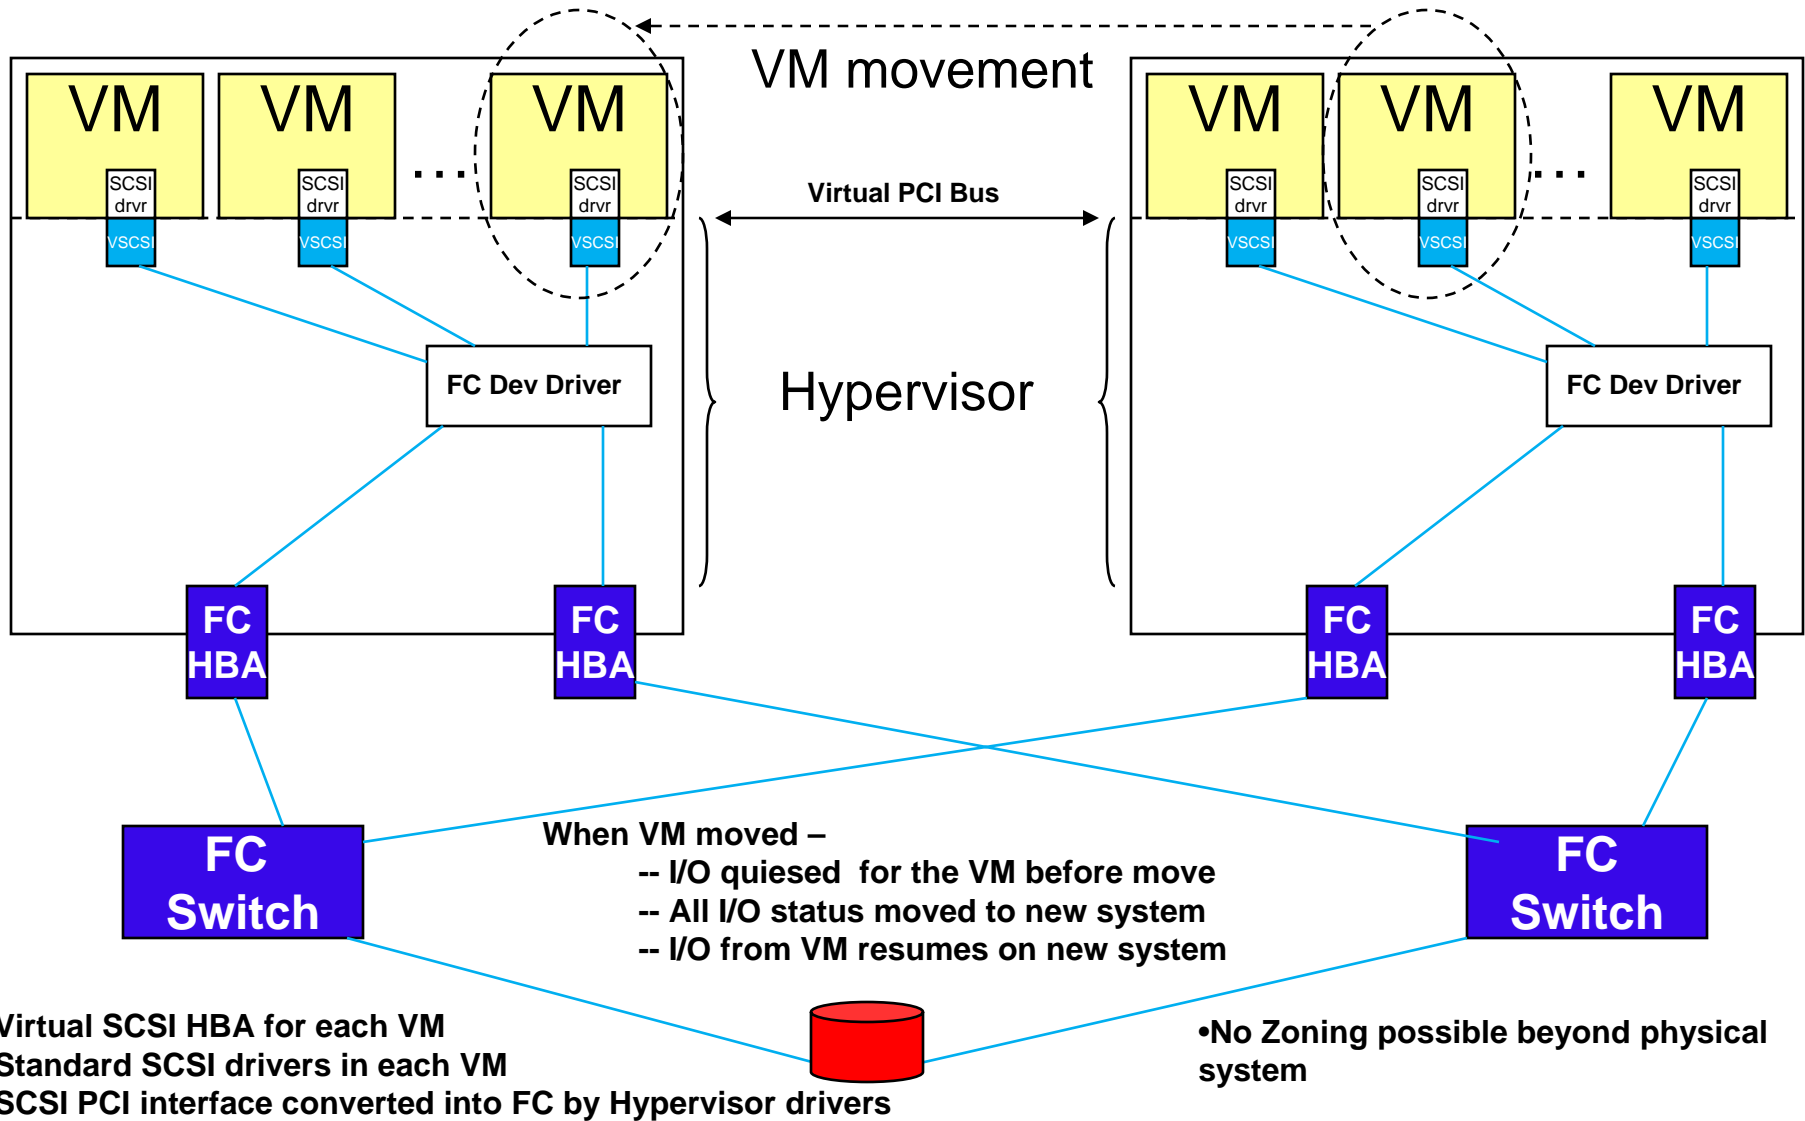


Current VMware with Vmotion (Ethernet)

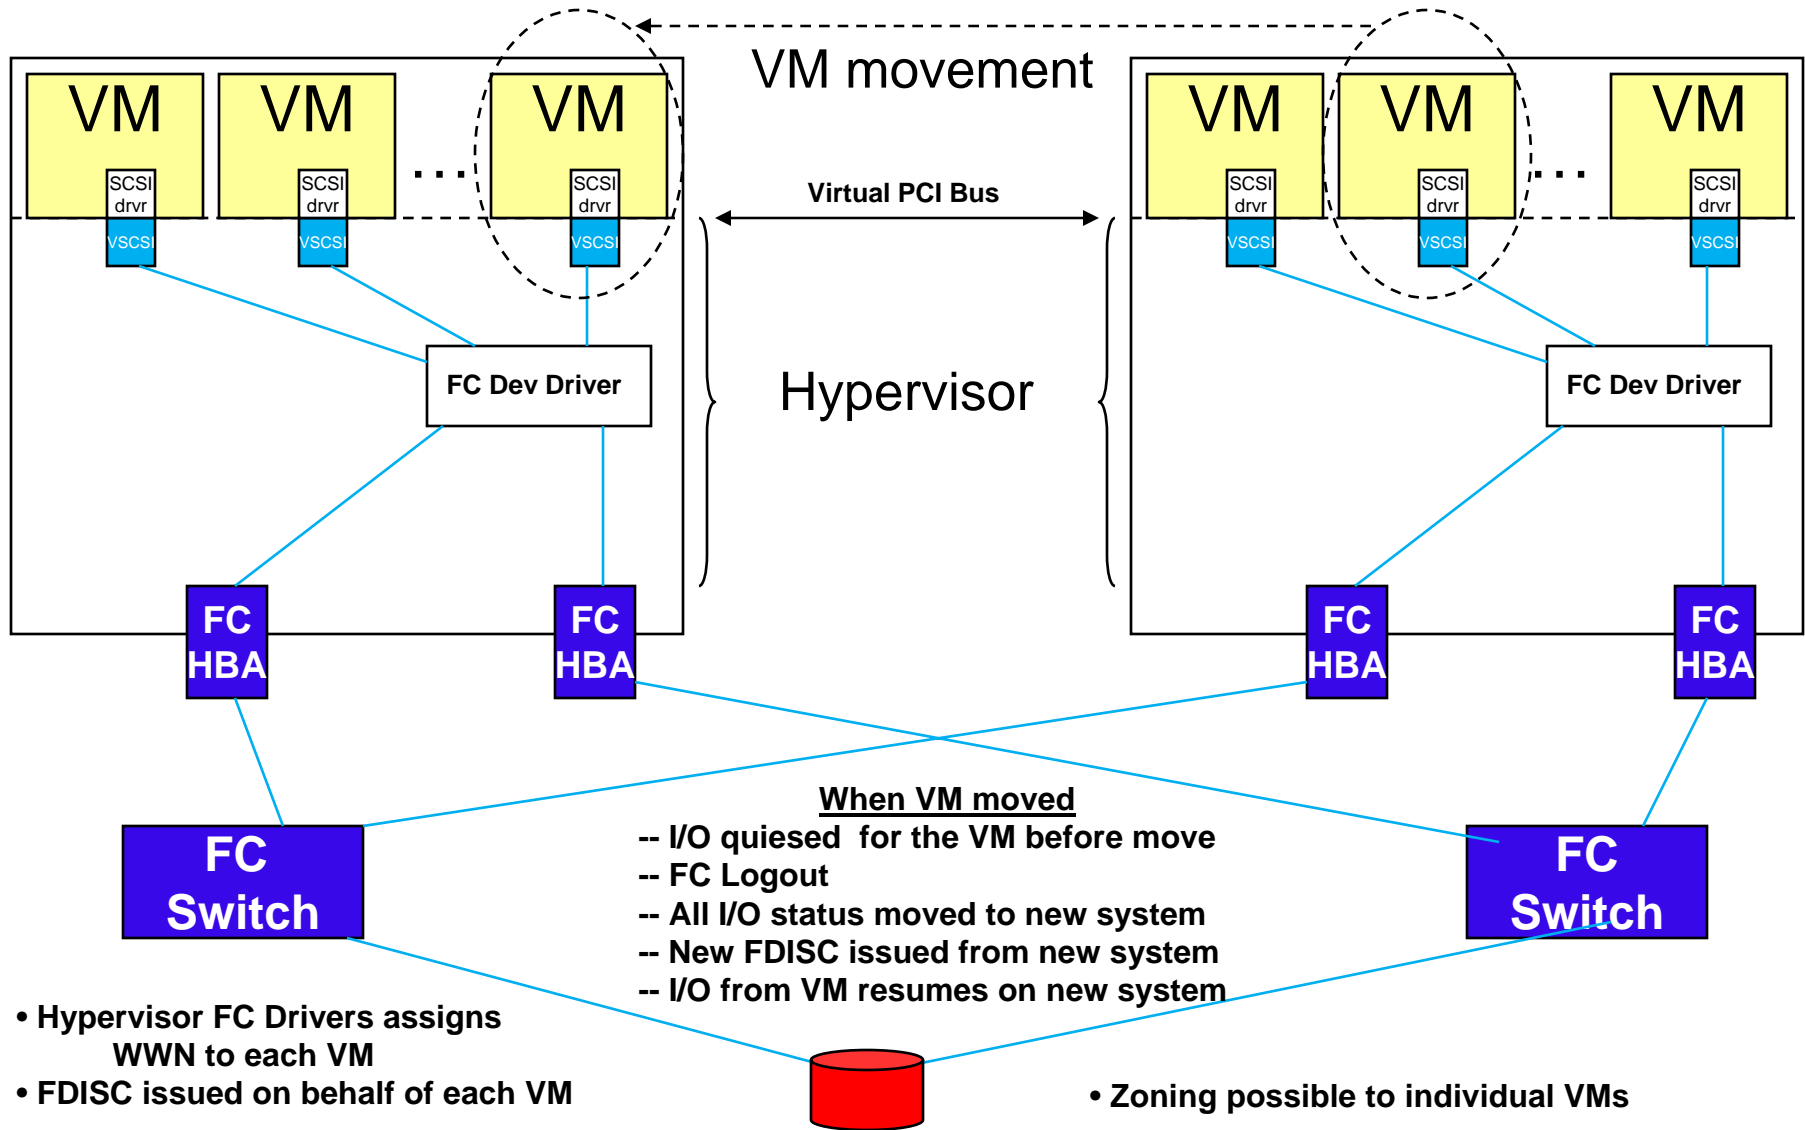


- Virtual NIC for each VM
- Standard NIC Drivers in each VM
- Virtual Switch within each Hypervisor
- Generates its own MAC addresses for each VNIC, MAC address moves with VM
- At completion of move, Ping from new system will notify the switches of new location of MAC


Current VMware with Vmotion (FC based RDM)



New VMware with Vmotion (NPIV FC based RDM)



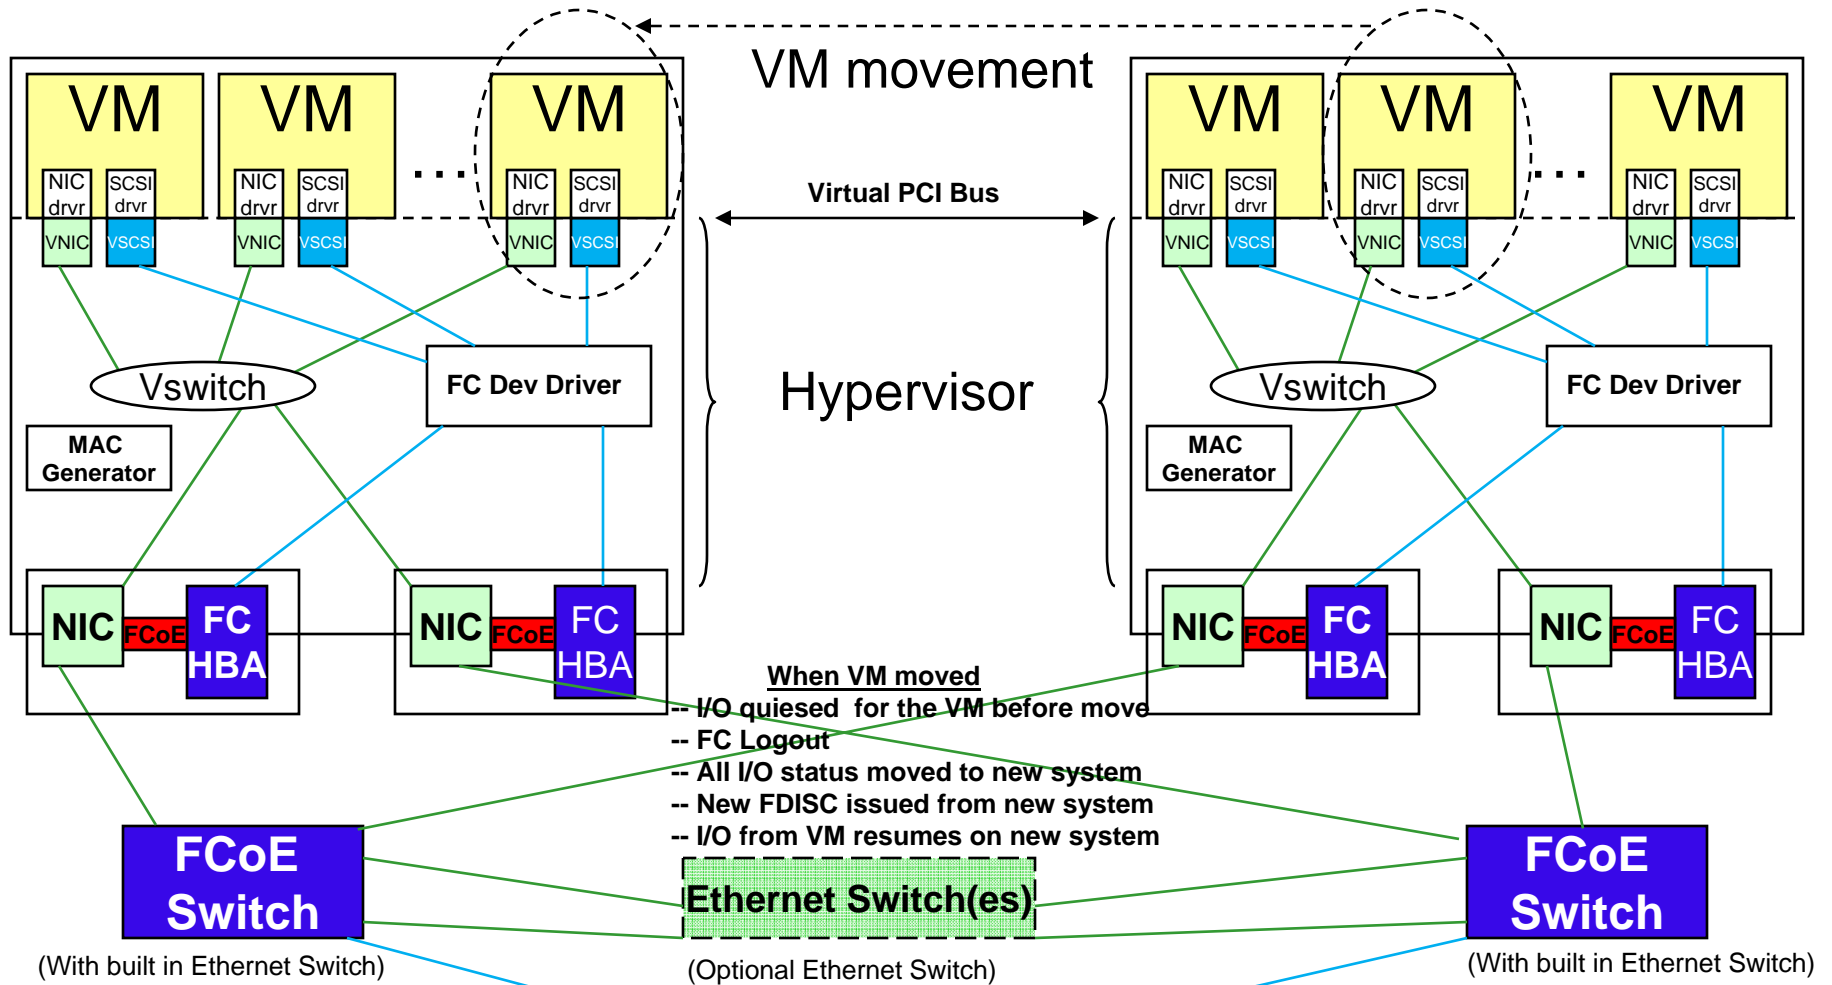
VMware with Vmotion (Ethernet and FC)



Vmotion & FCoE HBA Environments



VMware with Vmotion with FCoE/NIC Adapters



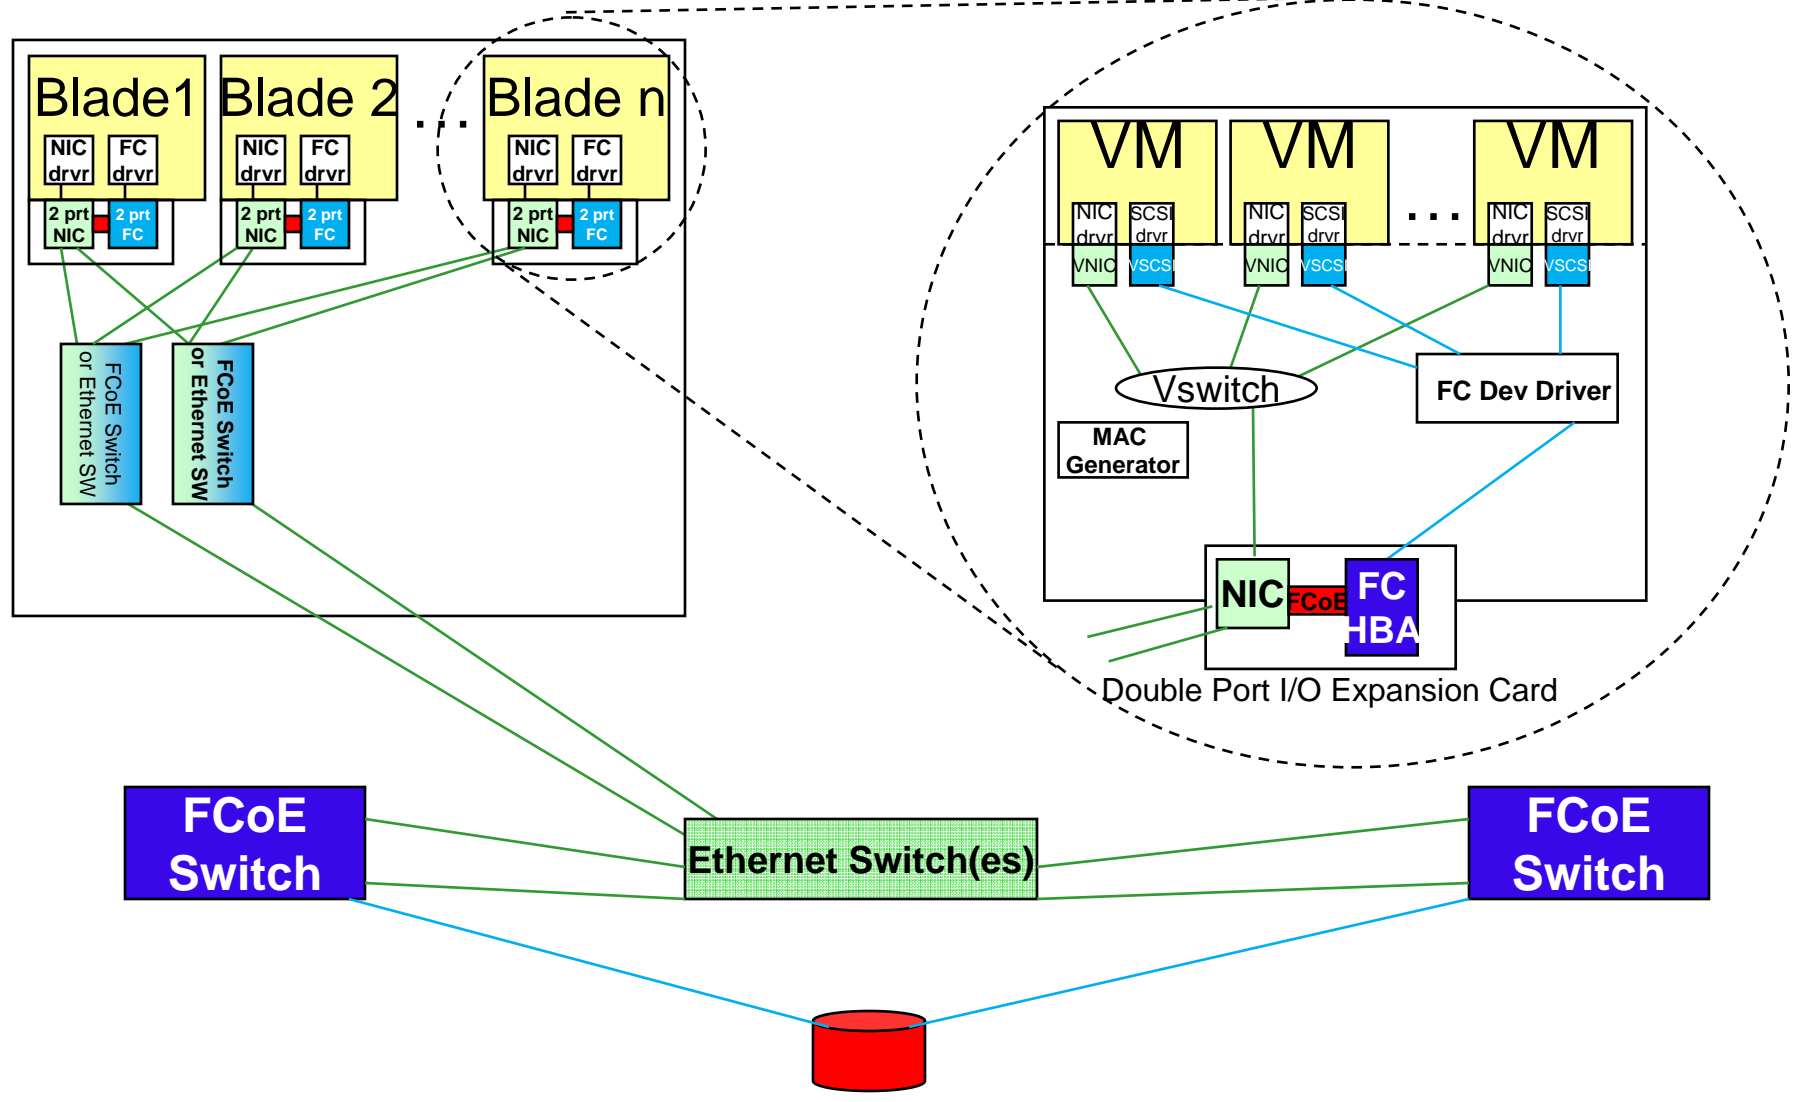
HBA has a FC compatible interface

Uses "Burnt-in" MAC Address

Same configuration for both NON NPIV and NPIV support



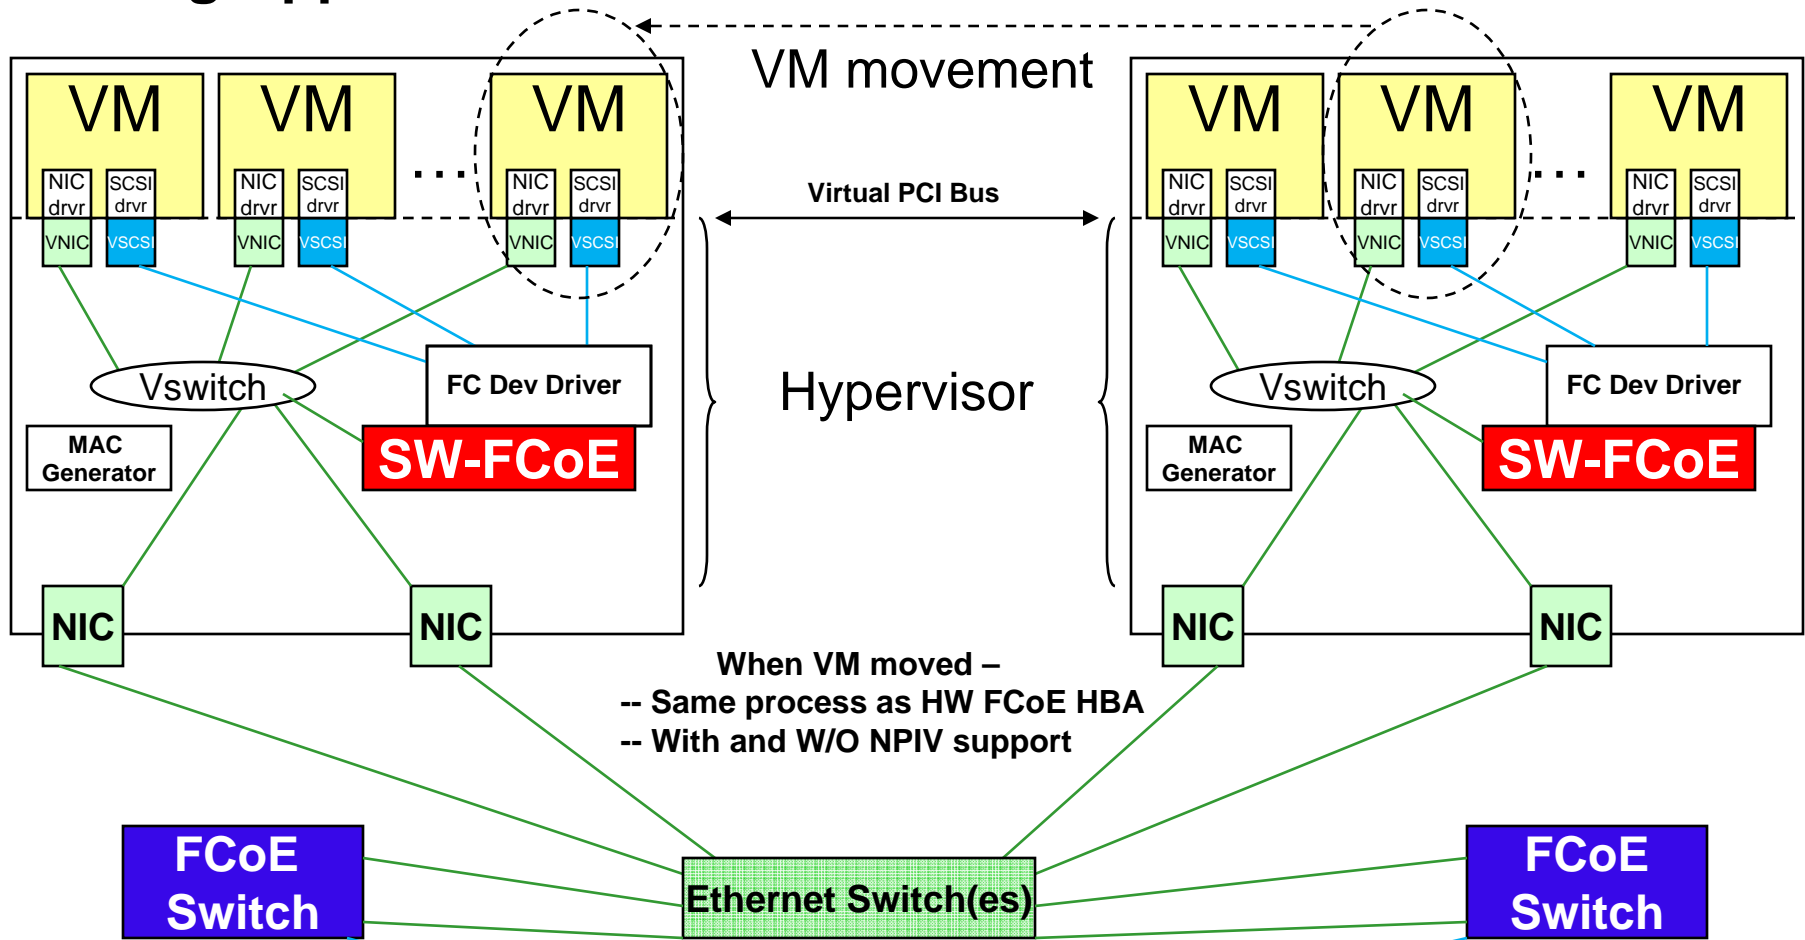
Blade Server with HW Virtualization on each Blade and FCoE Combo Switches (or Ethernet Switches)



Vmotion & Software FCoE Environments



Possible VMware with Vmotion & Software FCoE Using Approach 1



When VM moved –
 -- Same process as HW FCoE HBA
 -- With and W/O NPIV support

MAY use the Built-IN MAC Address
 SW-FCoE may interface to normal FC driver
 Same configuration for both NON NPIV and NPIV support

May also use MAC address
 obtained from the normal
 Hypervisor MAC Gen process
 (Server Provided MAC Addressing)



Notes regarding Software FCoE in the Hypervisor with Approach 1

The needed SINGLE MAC address for a Software FCoE implementation can be obtained from

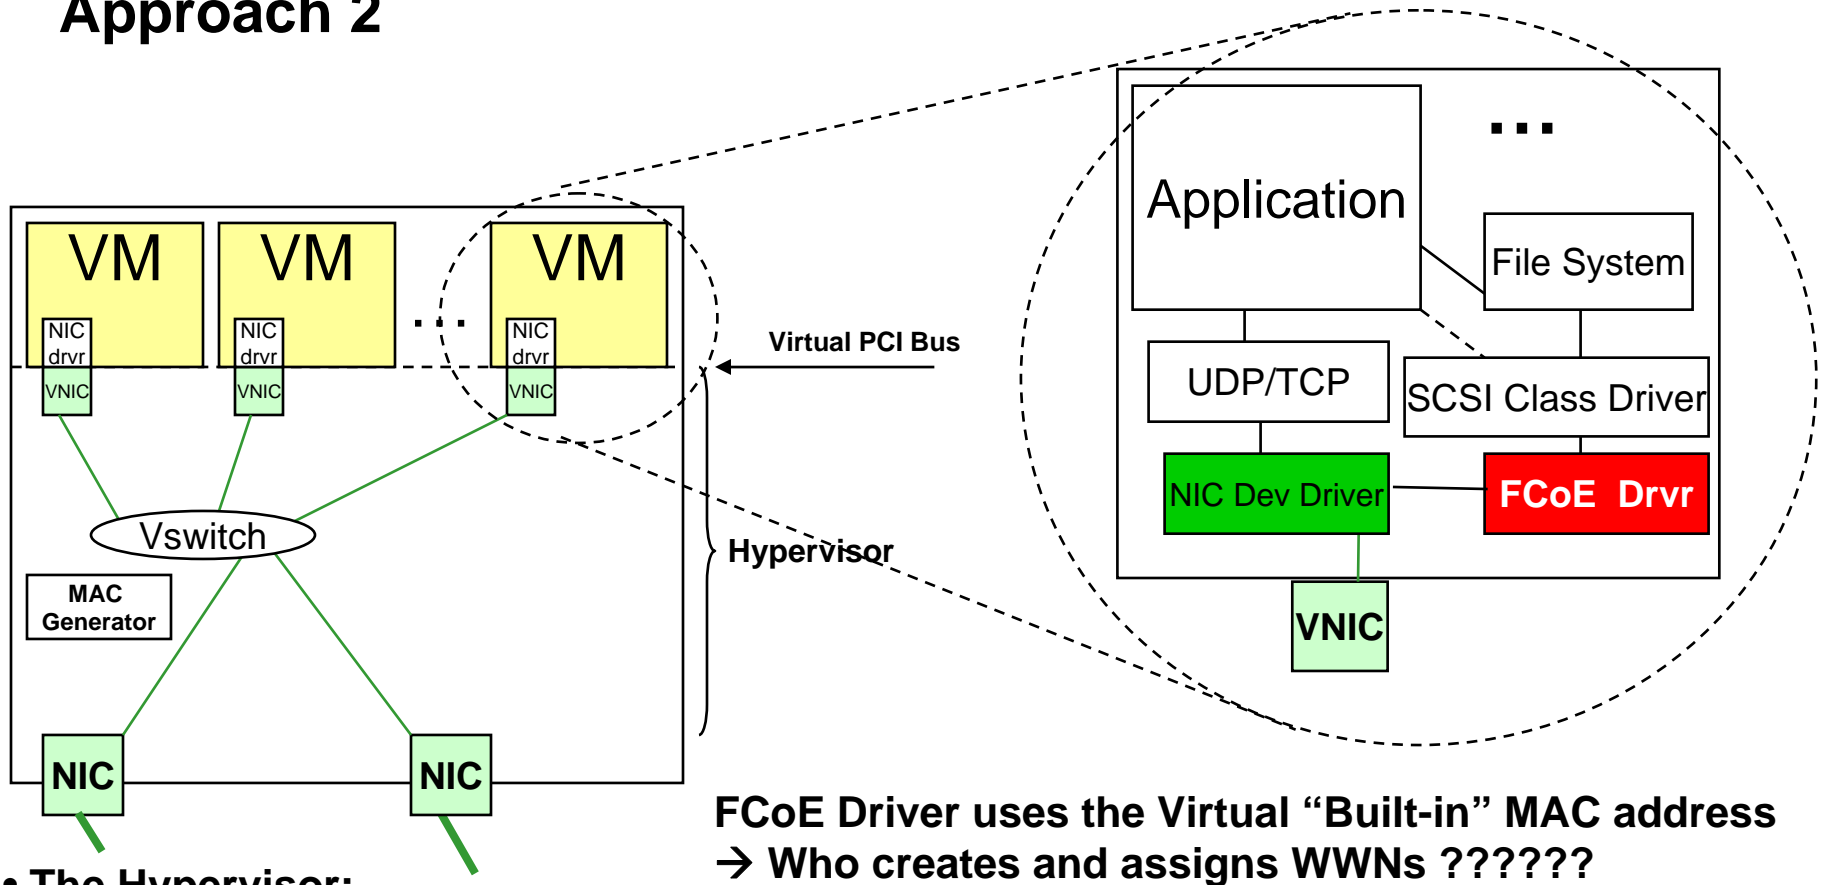
- A “Burnt-in” MAC address
 - Thus providing the same characteristics as Real FC
- or
- Via the normal Hypervisor/Admin MAC Generation process

→ It is all about the Server Providing the MAC address the way it does today

“Server Provided MAC Addressing”



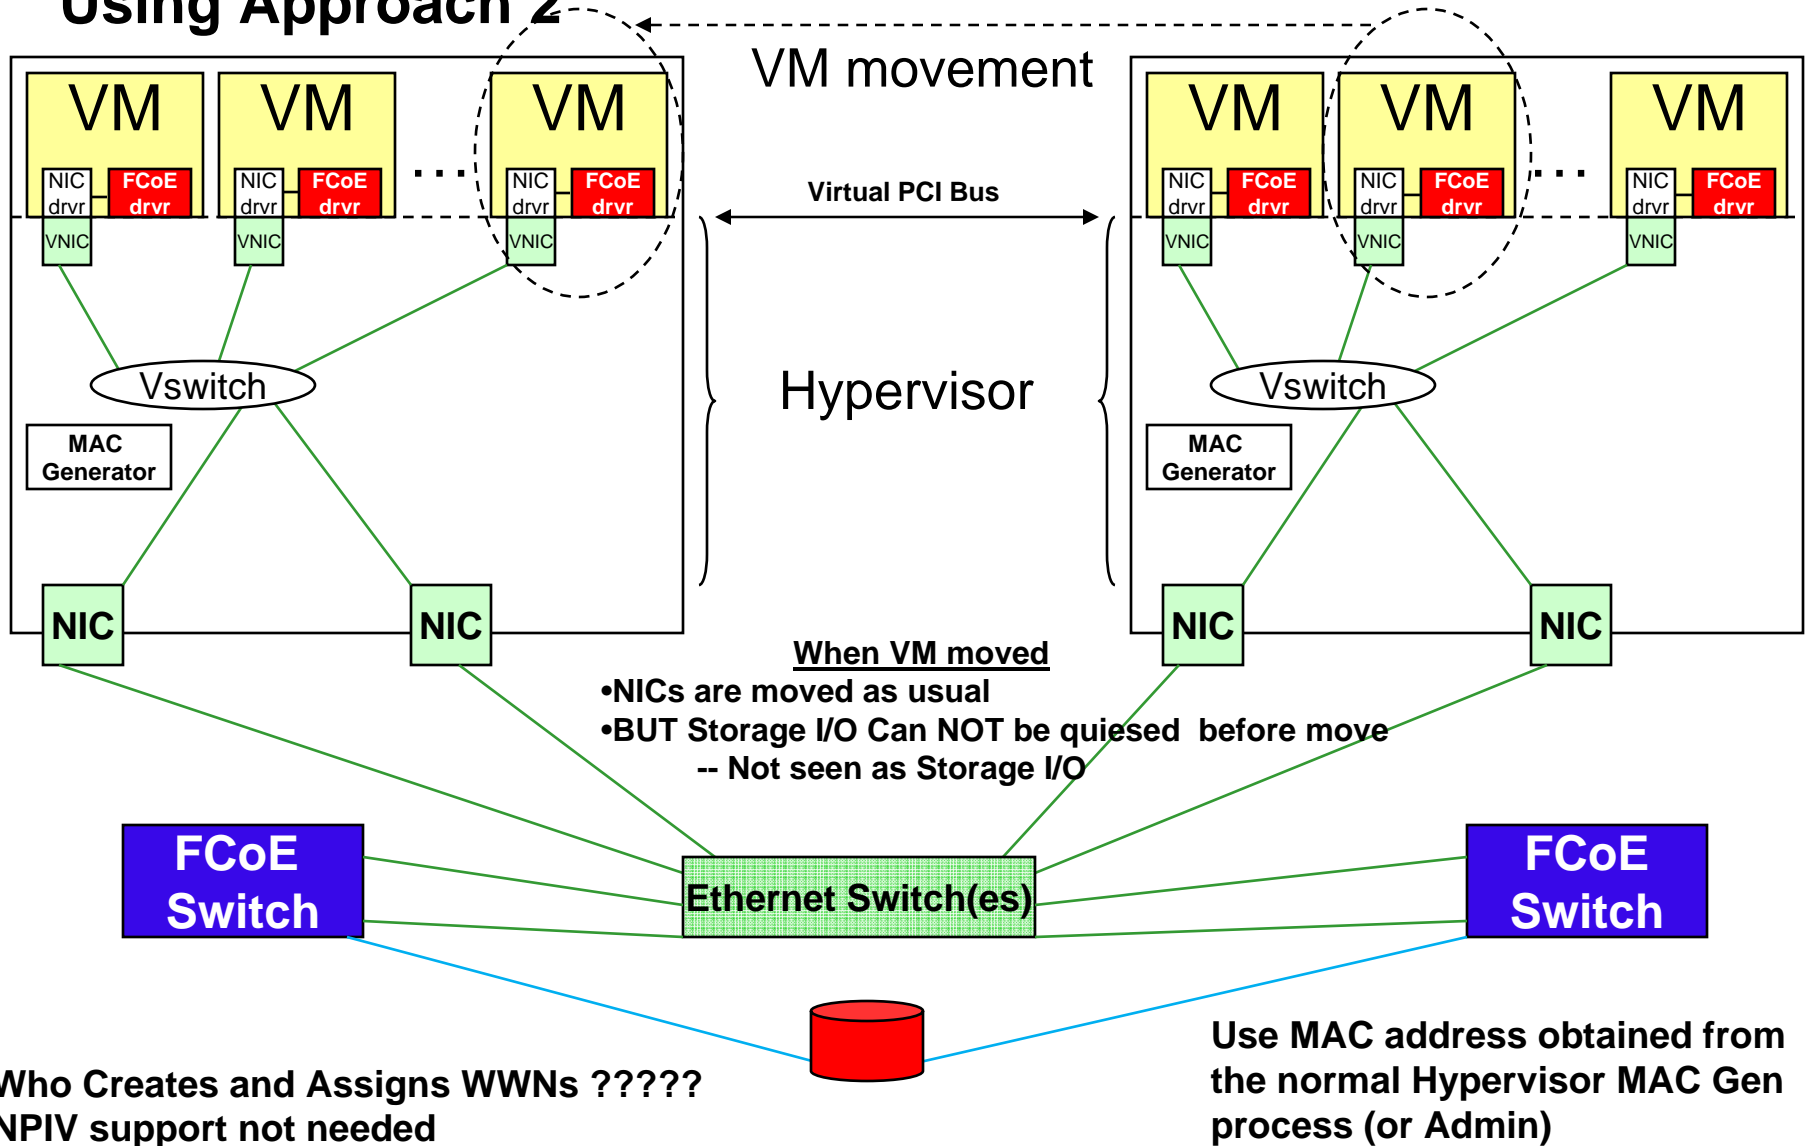
Remember Hardware Virtualization (with Software FCoE) Approach 2



**FCoE Driver uses the Virtual "Built-in" MAC address
→ Who creates and assigns WWNs ???????**

- **The Hypervisor:**
 - **Treats Storage I/O just like Network I/O**
(may be a problem with Vmotion – see later slides)
 - **Assigns the Virtual NIC MAC Address as usual**
(may have a separate VNIC for Storage)
- **NPIV support not needed**

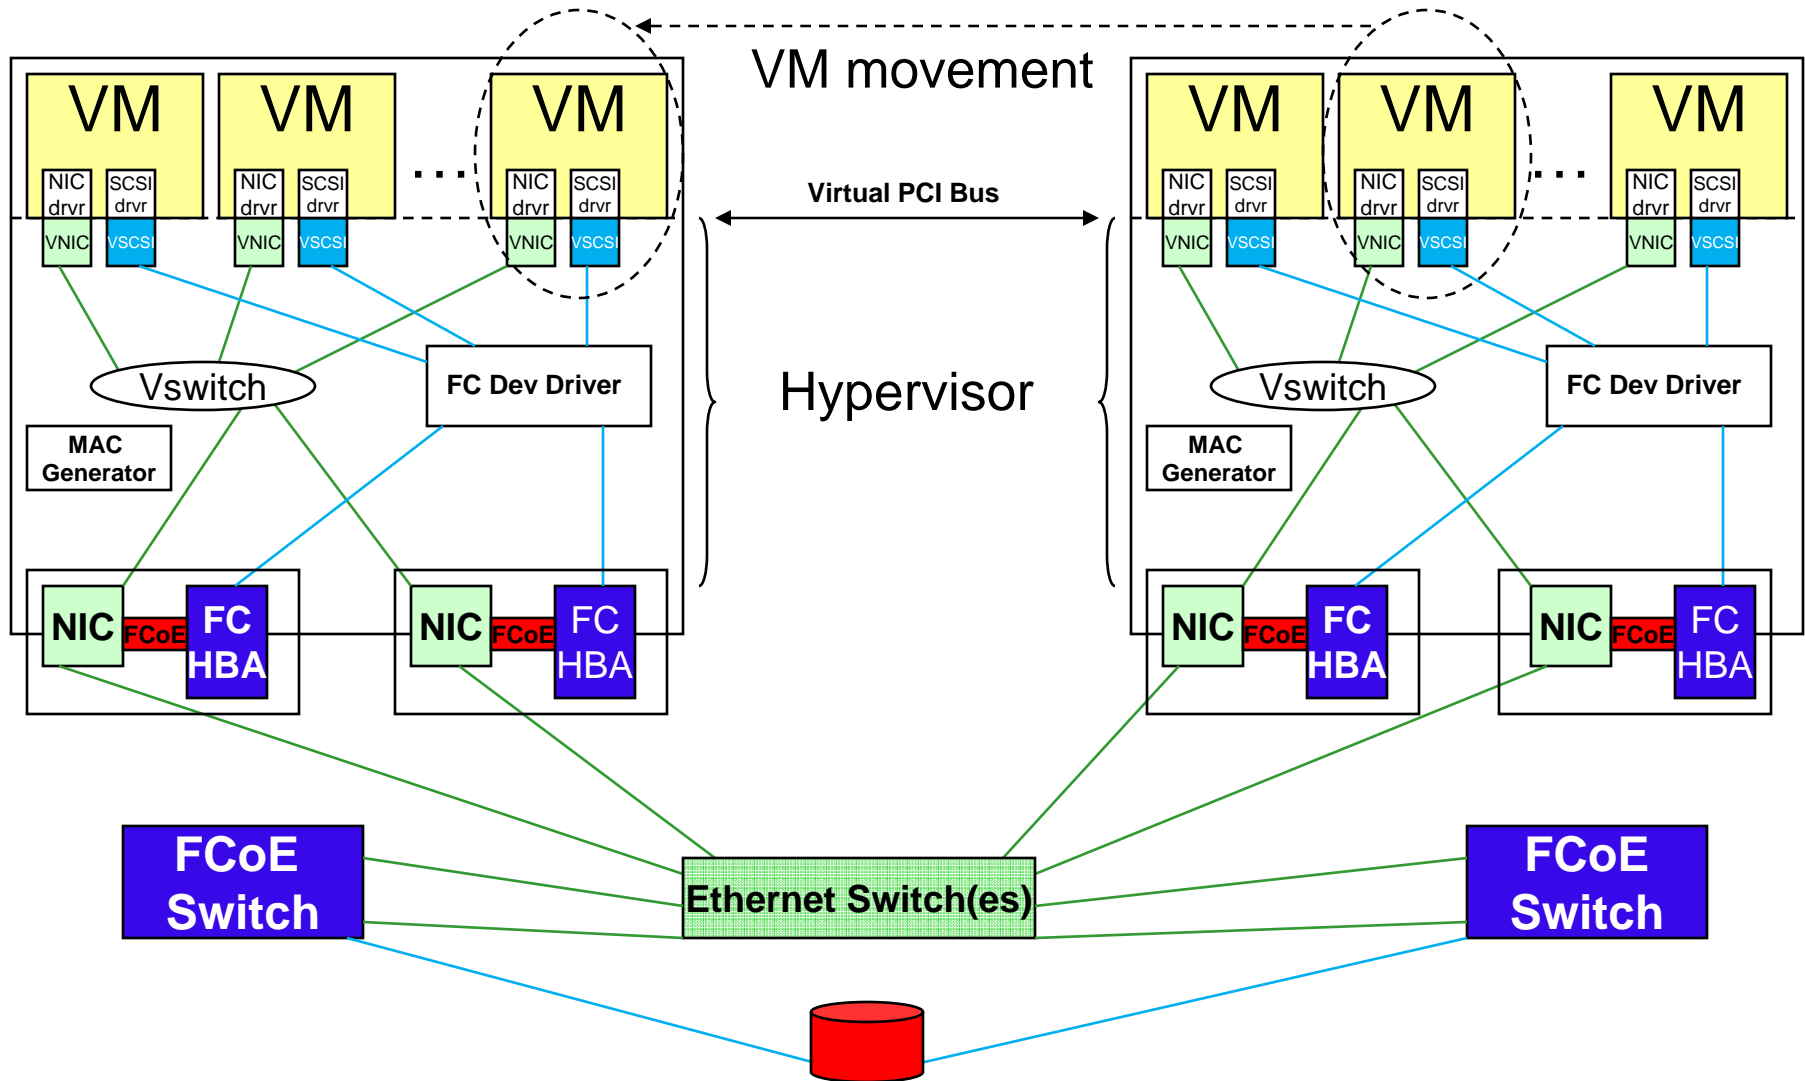
Possible VMware with Vmotion & Software FCoE Using Approach 2



Single Root I/O Virtualization (SRIOV)



VMware with Vmotion with SRIOV ????



Vendor Designs not yet exposed, but they need to be compatible with FC



Notes on SRIOV

Regardless of the implementation approach

- If the VMware SRIOV model continues to support FC HBAs
 - Then that support should also continue to support SRIOV FCoE HW adapters
 - SRIOV FCoE HW adapters can continue to use a Burnt-in MAC address for the FCoE operations



Summary

All major OSes currently support Real FC

- Took years to chase out all Network & HBA corner cases

New NPIV support is based on Real FC

- Will take lots of additional time to ensure that everything is addressed

Future SRIOV support will have to include Real FC

- Will take even more time to shake out all the implementation issues

We want to avoid special Driver Support for FCoE

- Want to take advantage of the drivers used for Real FC

All FCoE HW implementations that maintain a FC interface may use a “Burnt in” MAC Address

FCoE Hypervisor Software FCoE implementations may exploit “Server Provided MAC Addressing” by using:

- The Real NIC “Burnt-in” MAC address
- or
- A Hypervisor provided MAC address



Thank you



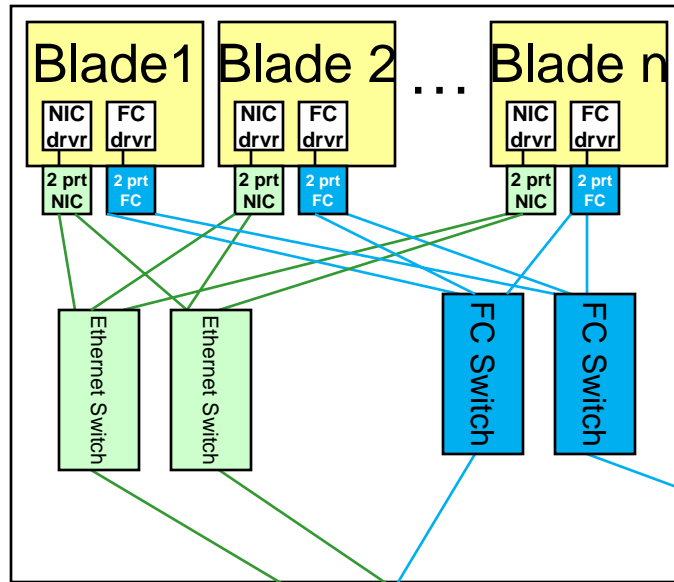
Backups Slides



Blade Servers




Blade Server with FC & NIC connections

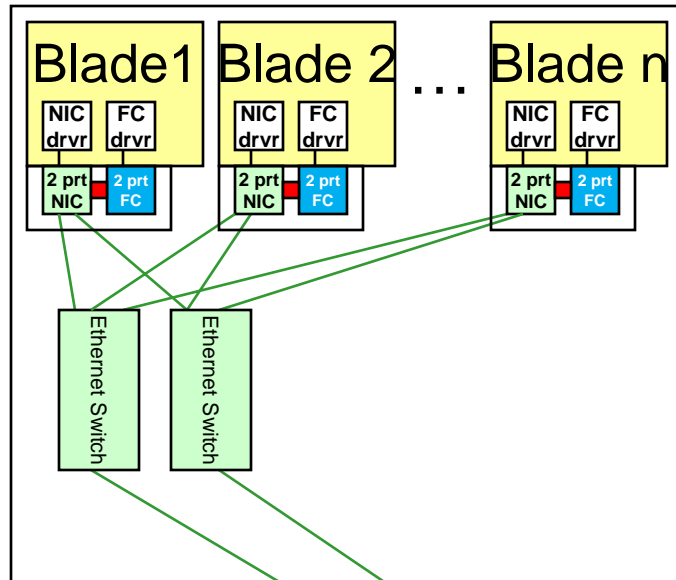


- Ethernet NICs or FC HBAs can be located on I/O expansion cards or on the Motherboard
- In this example, **each Blade has a single OS**
- Multiple connection types are needed FC, & Ethernet
- Redundant Switches needed for each connection type


Note: Outfacing NPIV implemented within the Blade Server FC switch is supported



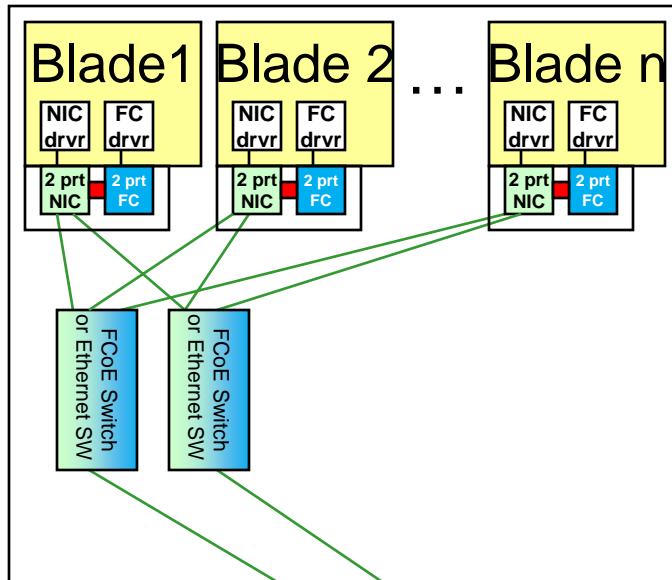
Blade Server with FCoE



- Ethernet FCoE HBAs can be located on I/O expansion cards (perhaps some day on the Motherboard)
- In this example, **each Blade has a single OS**
- Only Ethernet connections needed
- Only a Redundant Switch set is needed
- Only NIC “Burnt-in” MAC Addresses are needed
- Notes:
 - The physical configuration looks the same when Virtualizing OS is operating on the blades
 - No need for outfacing NPIV support needed in the imbedded Ethernet Switches



Blade Server with FCoE with Combo Switches



- Ethernet FCoE HBAs can be located on I/O expansion cards (perhaps some day on the Motherboard)
- In this example, **each Blade has a single OS**
- Only Ethernet connections needed
- Only FCoE switches with integrated Ethernet switches
- Only NIC “Burnt-in” MAC Addresses are needed
- Outfacing NPIV support possible but probably not needed in the FCoE imbedded switches

