NPIV

FCP Channel Virtualization in a Linux Environment

Volker Sameske (sameske@de.ibm.com)
Linux on zSeries Development
IBM Lab Boeblingen, Germany

Share Baltimore, Maryland
August 13-18, 2006
Session 9257

ON DEMAND BUSINESS™

FCP Channel Virtualization

- What is NPIV?
- New Possibilities
- Requirements
- Linux Support
- Getting Started
- Look 'n Feel

Terms and Definitions

- NPIV
  - N-Port ID Virtualization
  - IBM System z9 only!
- Channel
  - Physical FCP adapter
- Subchannel
  - Logical FCP adapter
  - Defined within IOCDS or z/VM
- HBA
  - Host Bus Adapter (channel or subchannel, context-dependent)
- FC SAN
  - Fiber Channel Storage Area Network

Terms and Definitions (cont.)

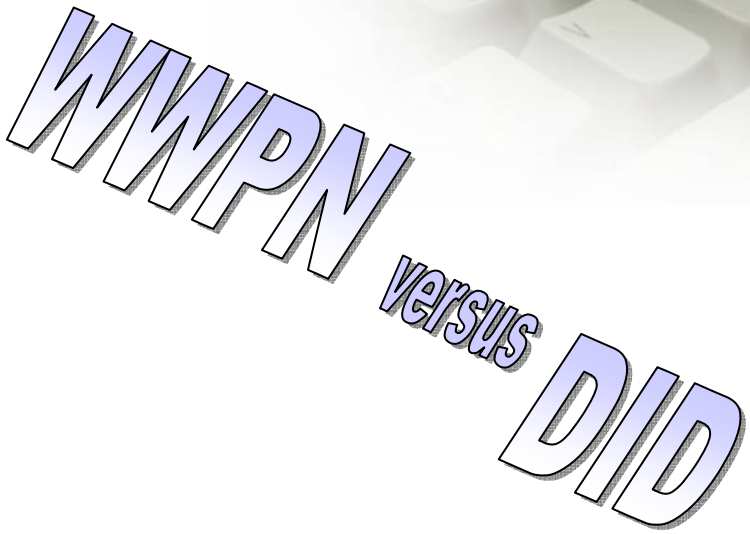

- WWPN
  - Worldwide Port Name
  - Worldwide unique identifier of a FC port
  - „burned in“ during manufacturing
  - 8 byte hexadecimal value
  - E.g. 0x5005076401e06b18
- D_ID
  - Destination ID
  - Used internally for active FC connections instead of WWPNs
  - 3 byte hexadecimal value
  - Assigned by FC switch
  - E.g. 0x2a036bc

WWPNs and other Ports

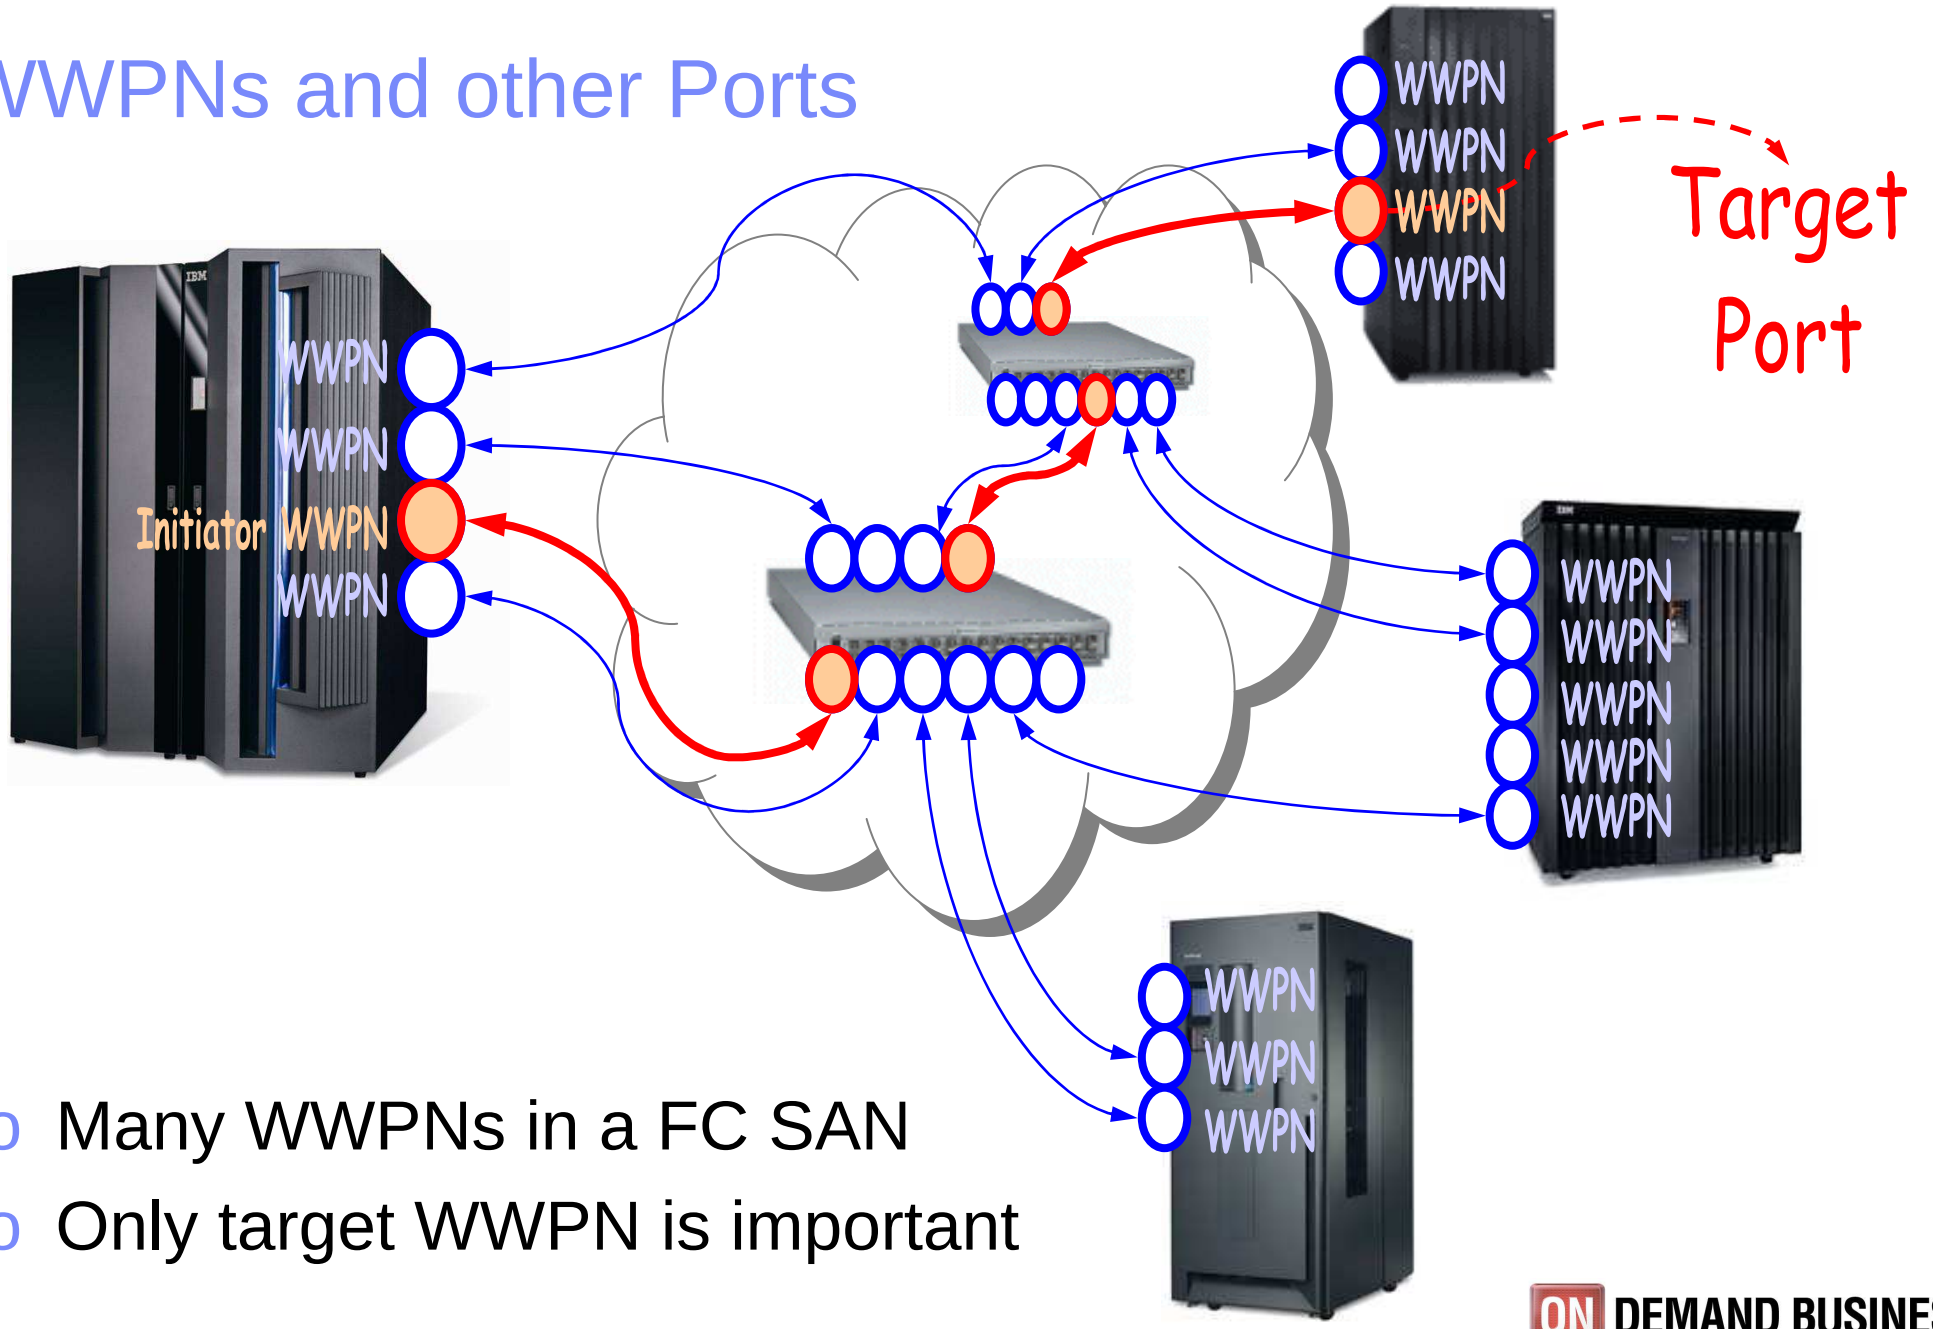

- Many WWPNs in a FC SAN
- Only target WWPN is important

What is NPIV?

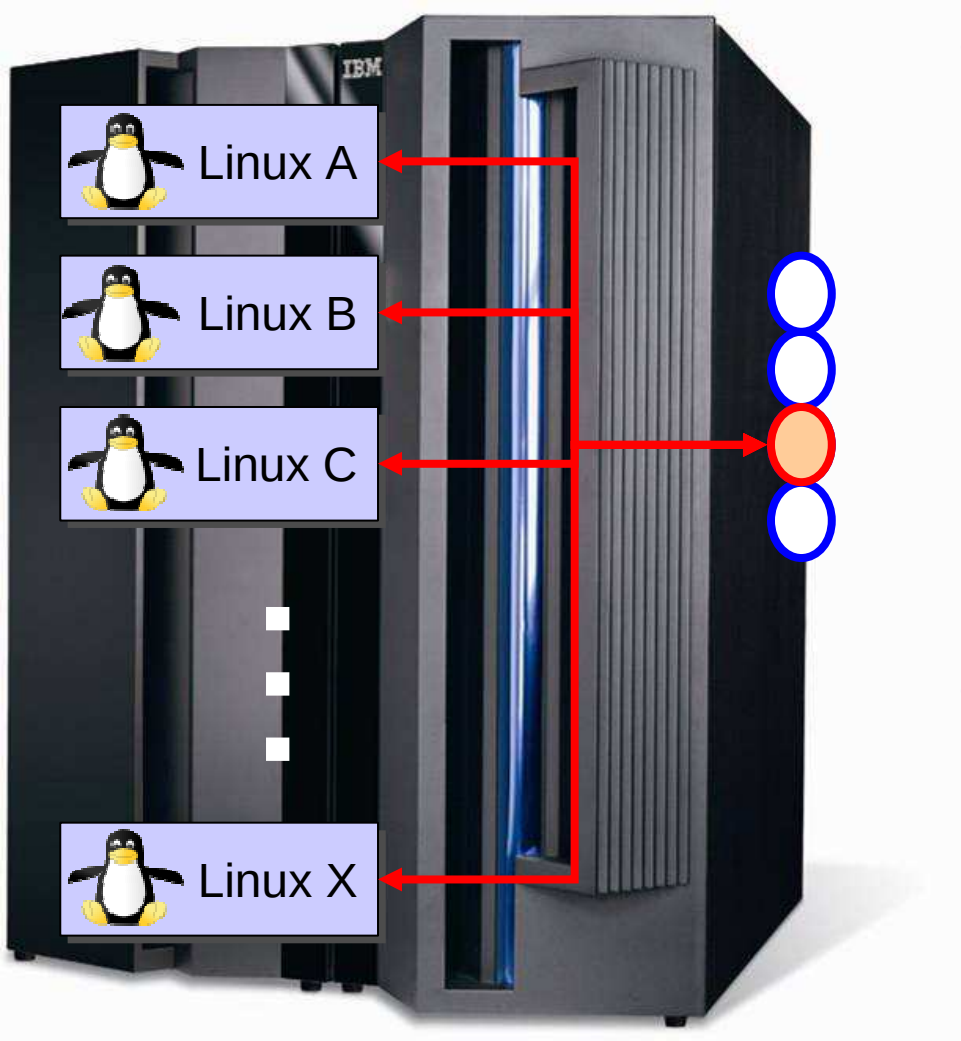

- Initiator Port
  - has attributes like:
    - WWPN
    - D_ID
    - PCHID
    - CHPID
    - …
- Without NPIV: All channel attributes will be duplicated for each subchannel

What is NPIV?

"Think Virtual Adapters"

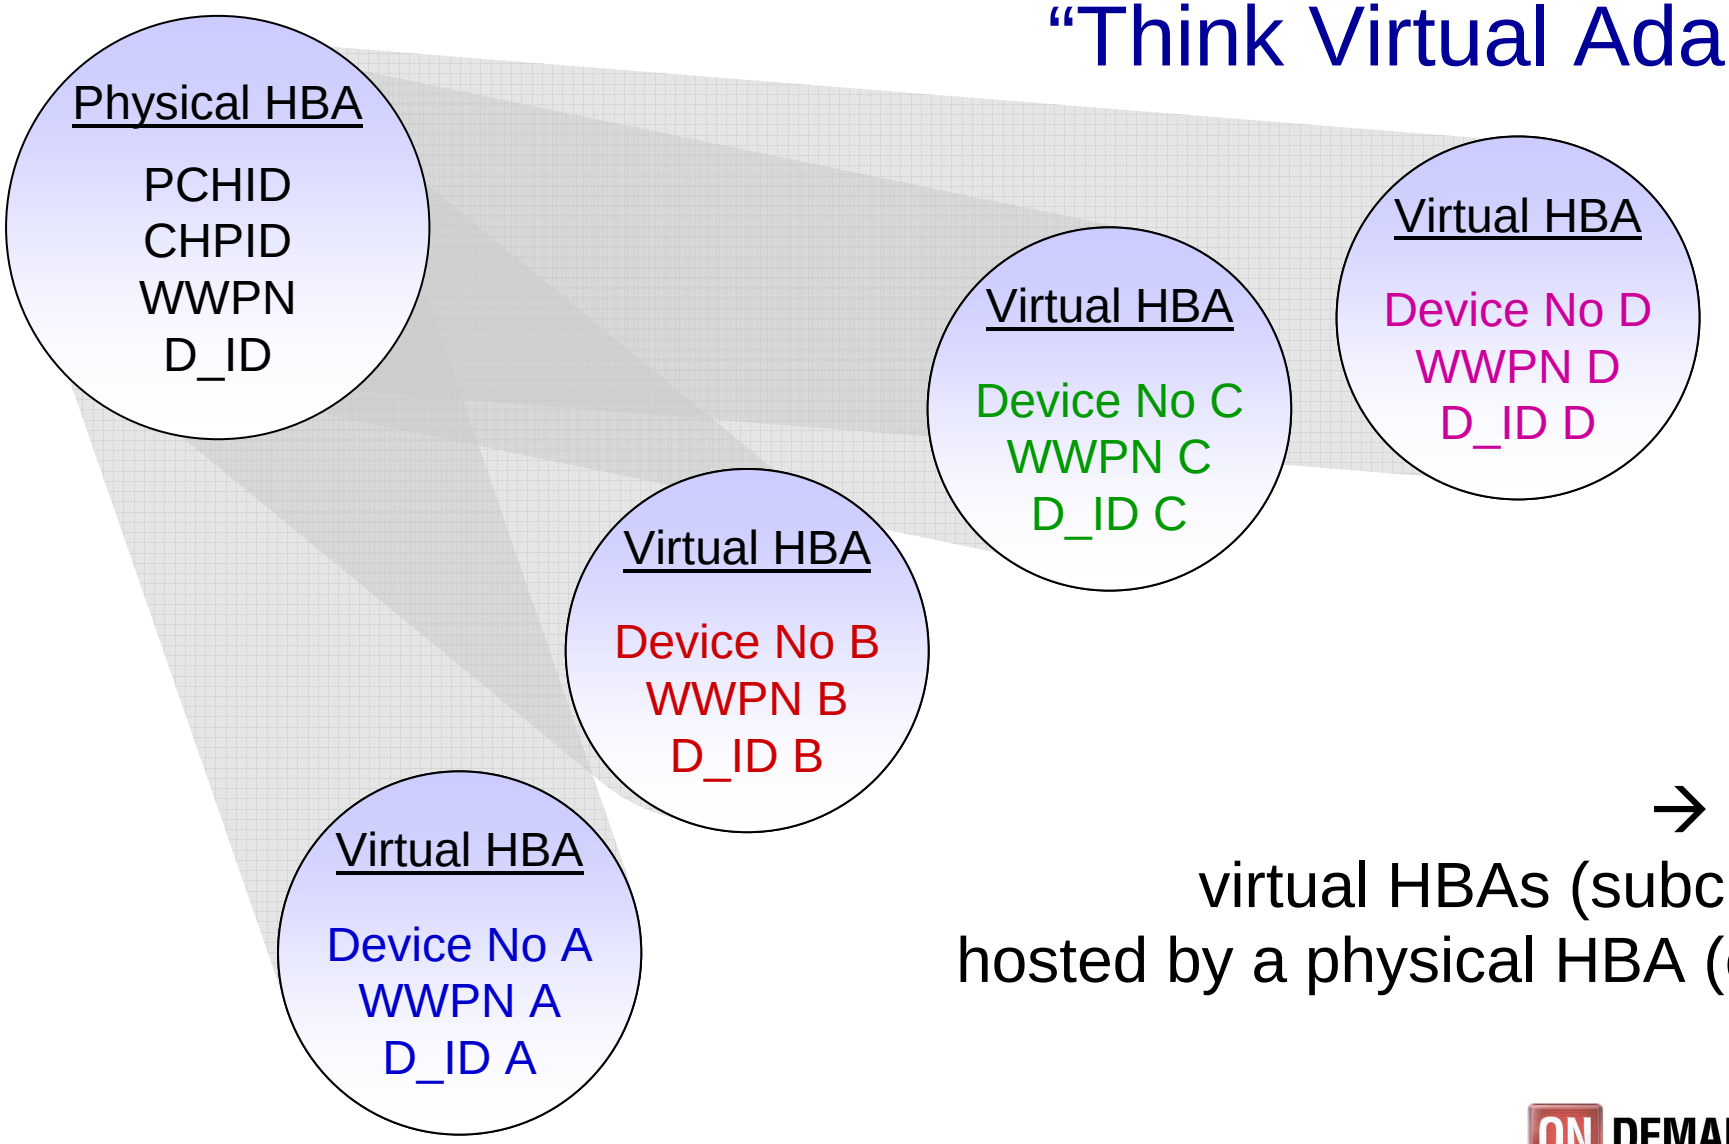

→ concept:
virtual HBAs (subchannels)
hosted by a physical HBA (channel)

What is NPIV?

Unique SAN Identities!

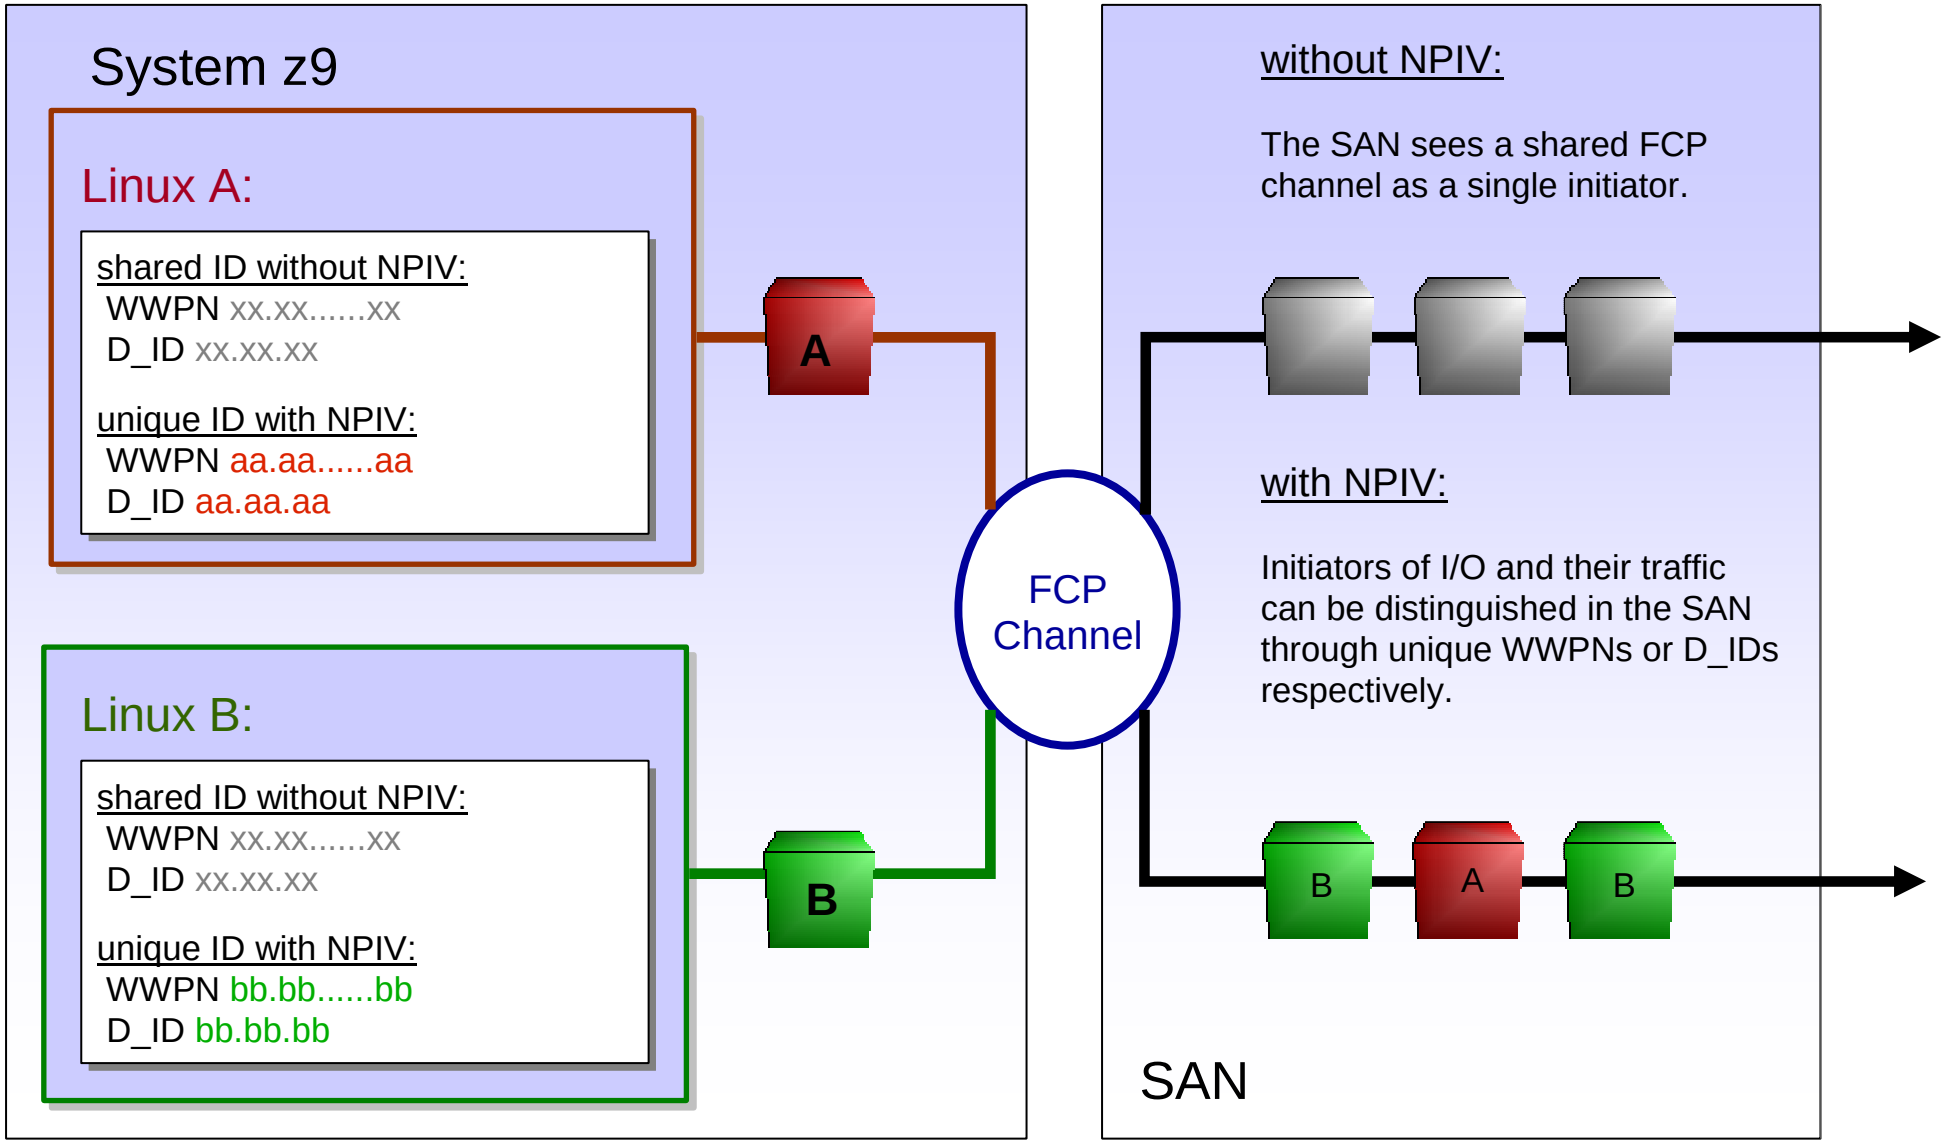

What is NPIV?

Industry-Standard Solution

- NPIV = N-Port Identifier Virtualization
- Standard-based approach being embraced by the industry
- System z9 persistently assigns unique WWPN to each FCP subchannel
- FCP Channel obtains separate D_ID for each subchannel from fabric switch
- result: unique SAN identity for each FCP subchannel

see also	Fiber Channel - Framing and Signaling Interface (FC-FS)	http://www.t11.org

What is NPIV?

System z9 Implementation

- Connecting a virtual HBA to the SAN

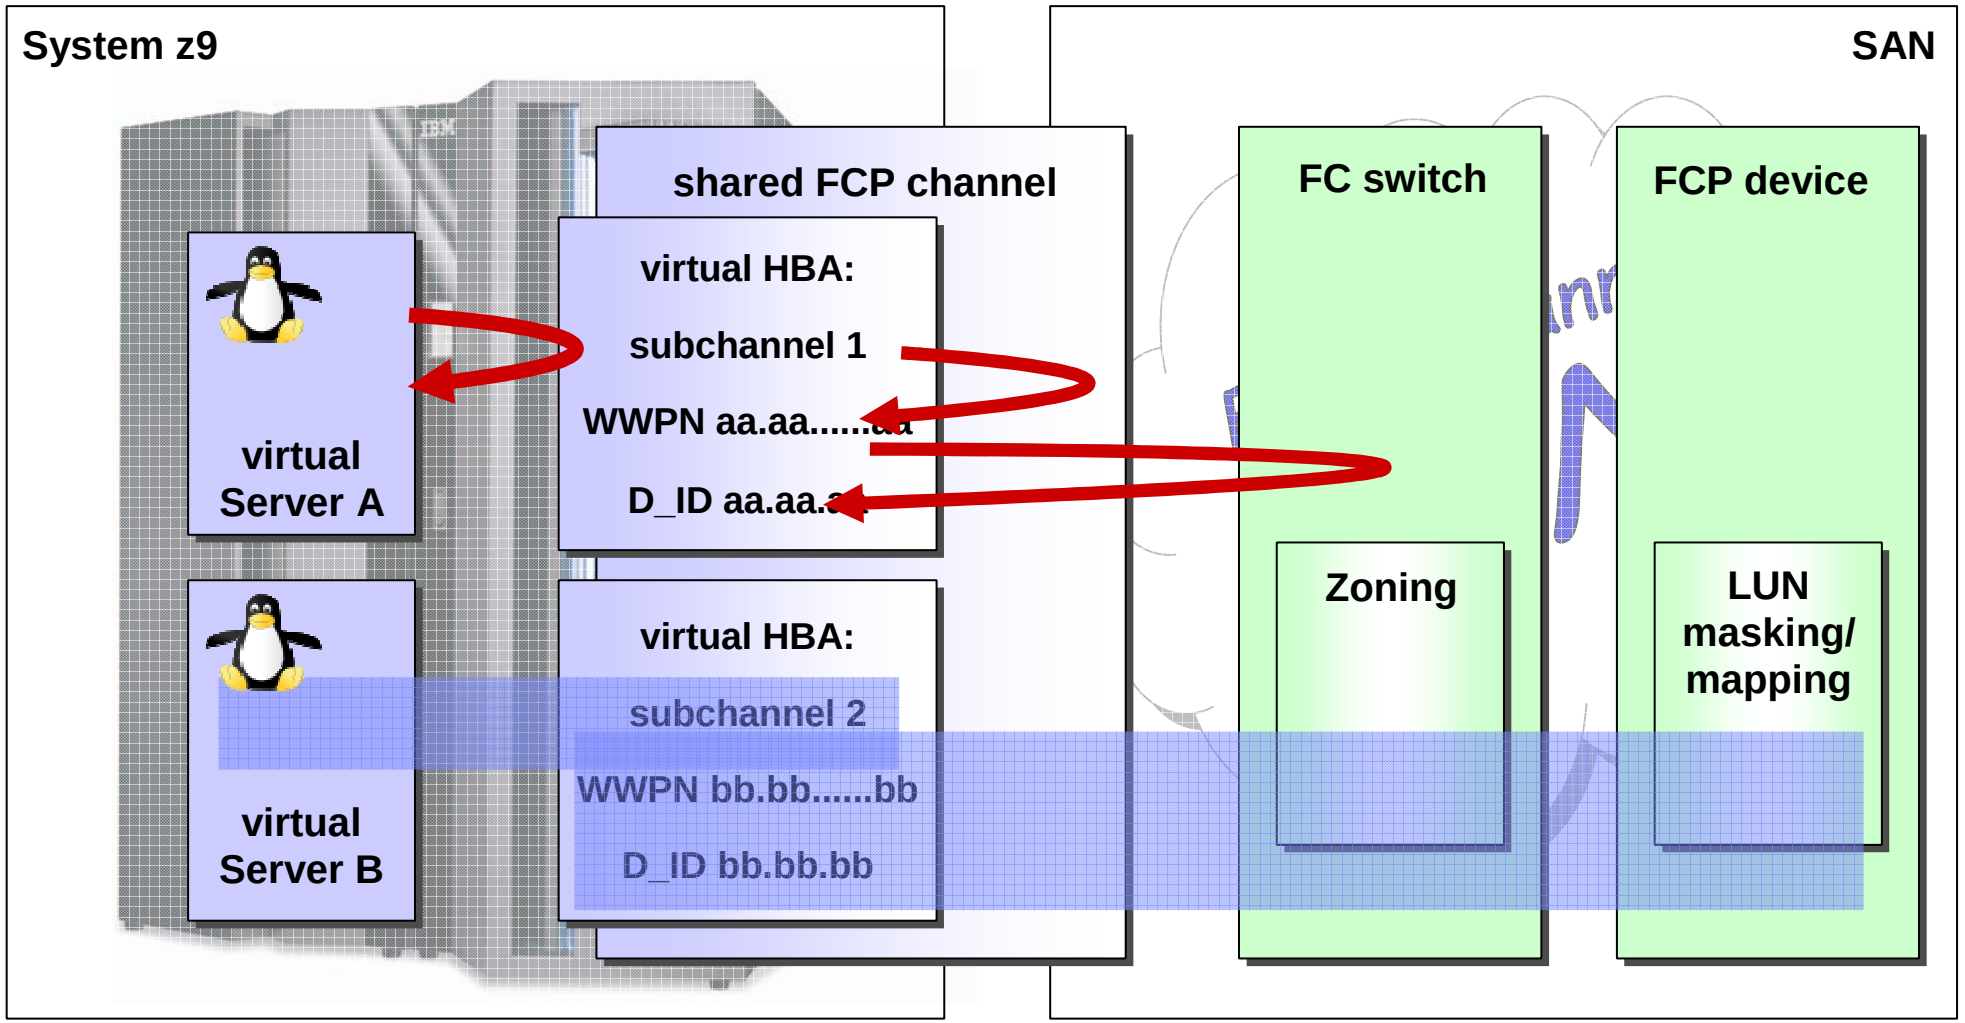

What is NPIV?

System z9 Implementation

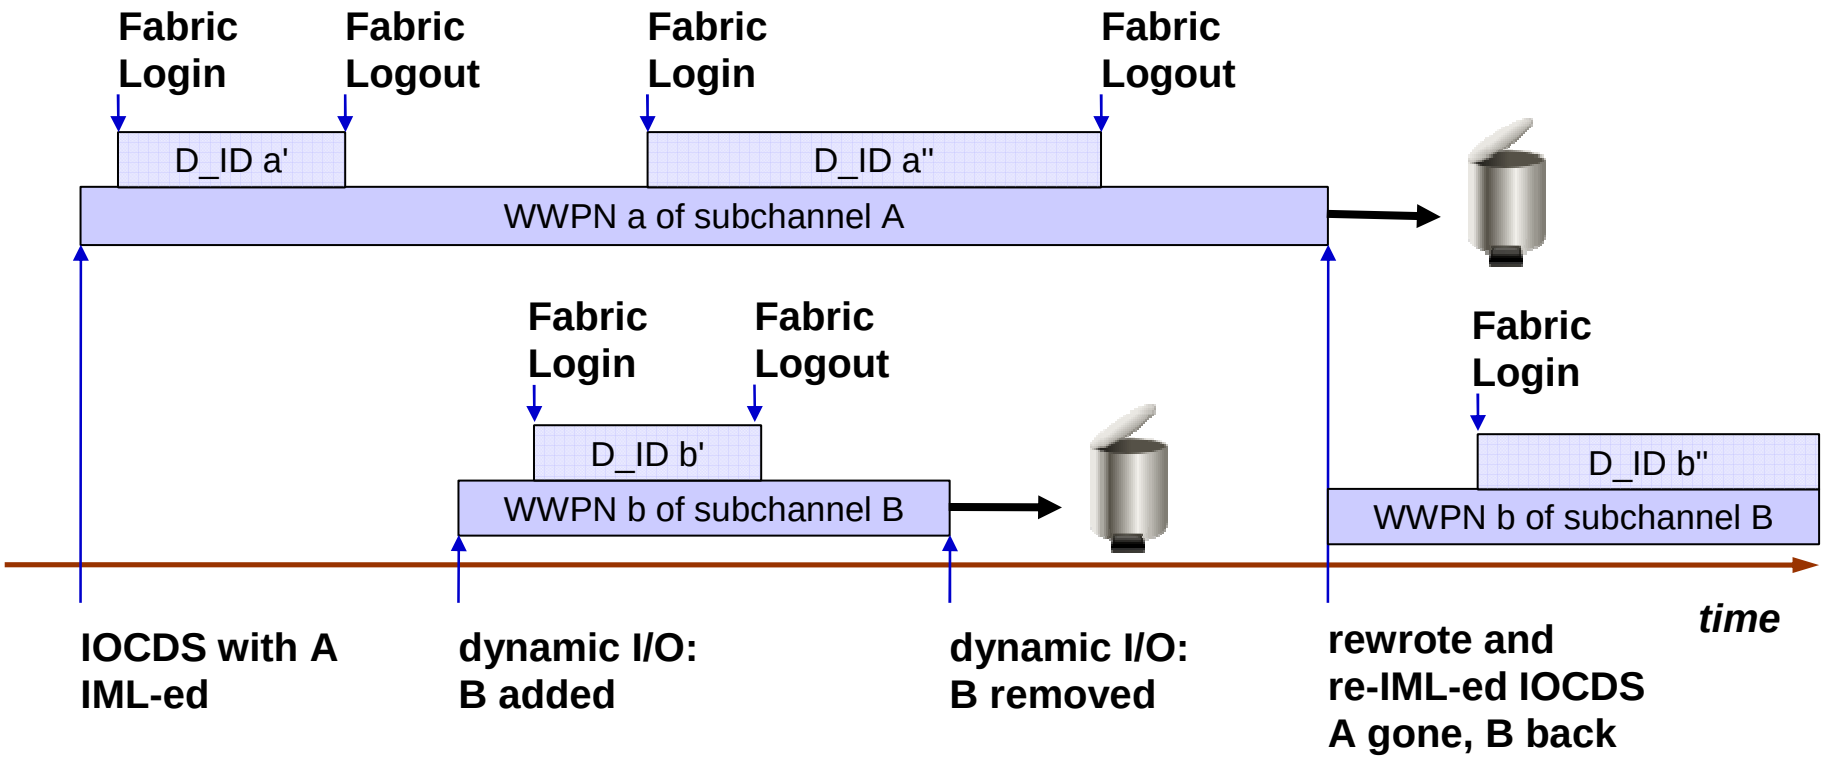

- Lifetime of virtual WWPN = lifetime of subchannel definition
- D_ID lifetime = fabric connection lifetime

What is NPIV?

System z9 Implementation

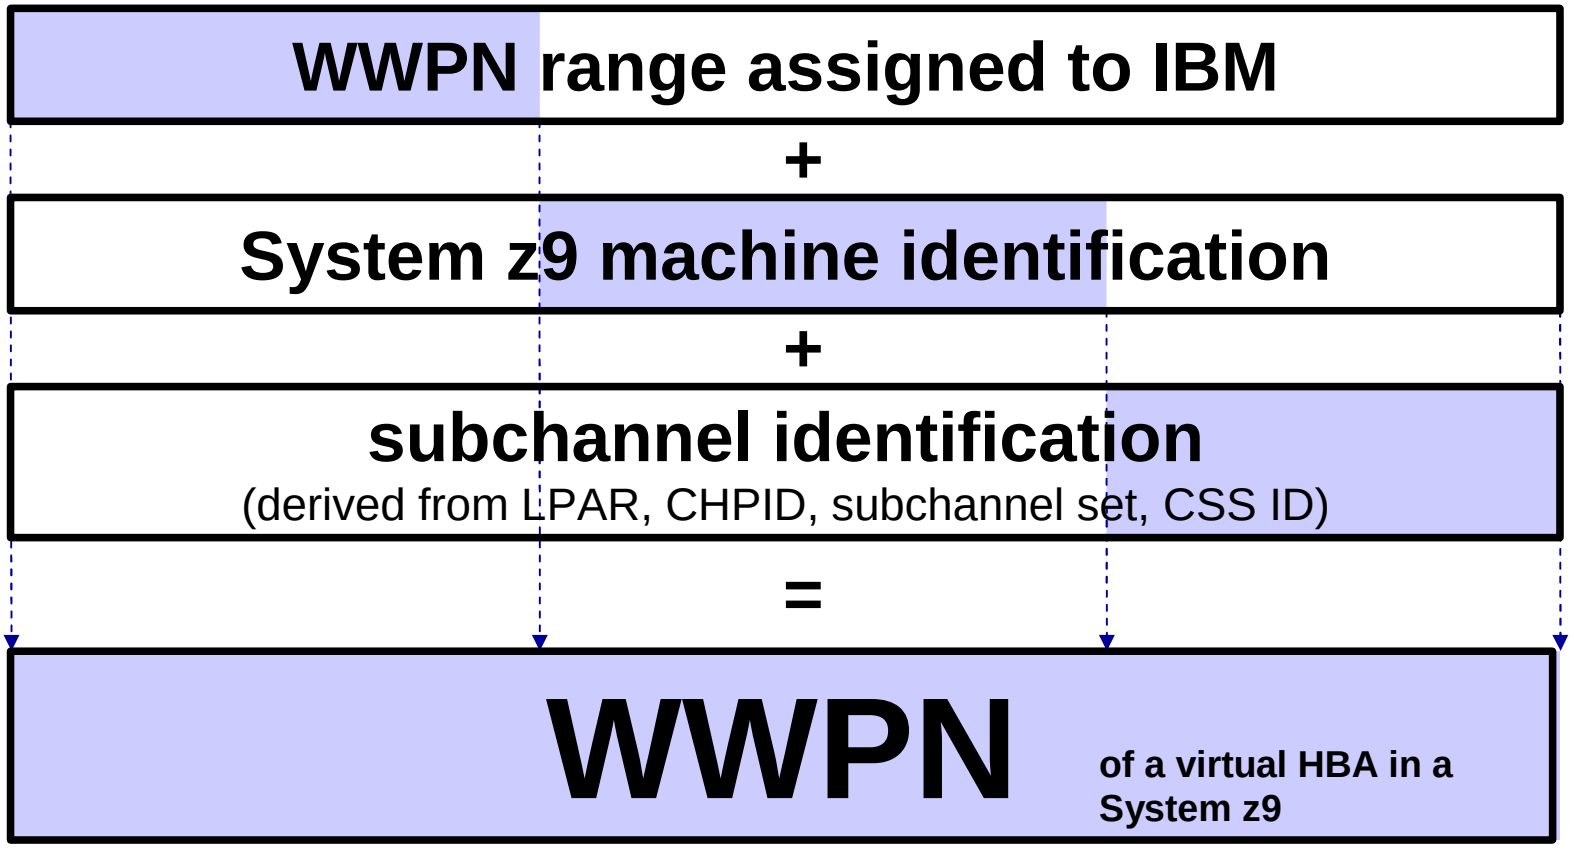

- well-defined “ingredients” tie a unique WWPN permanently to a subchannel

What is NPIV?

System z9 Implementation

- Migrating Linux changes WWPNs

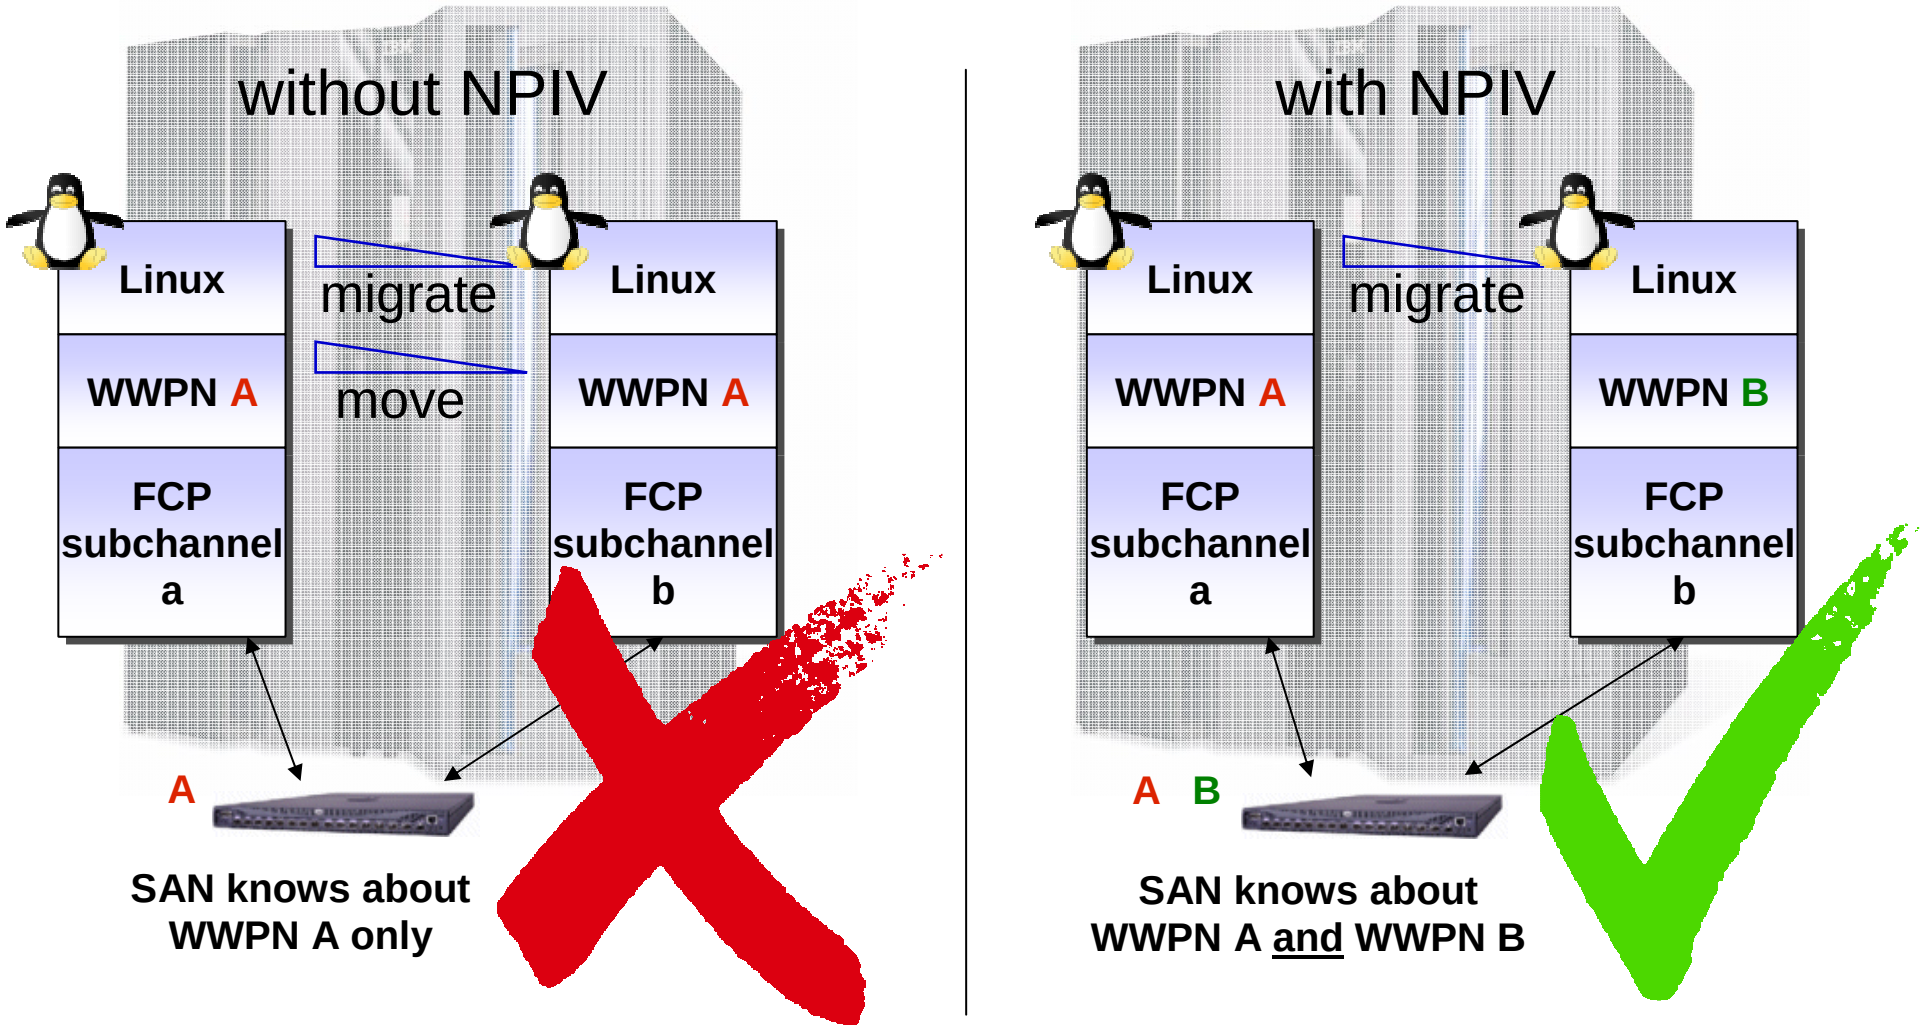

New Possibilities

'Exclusive LUN' policy eliminated

- Many-to-many tape backup solutions

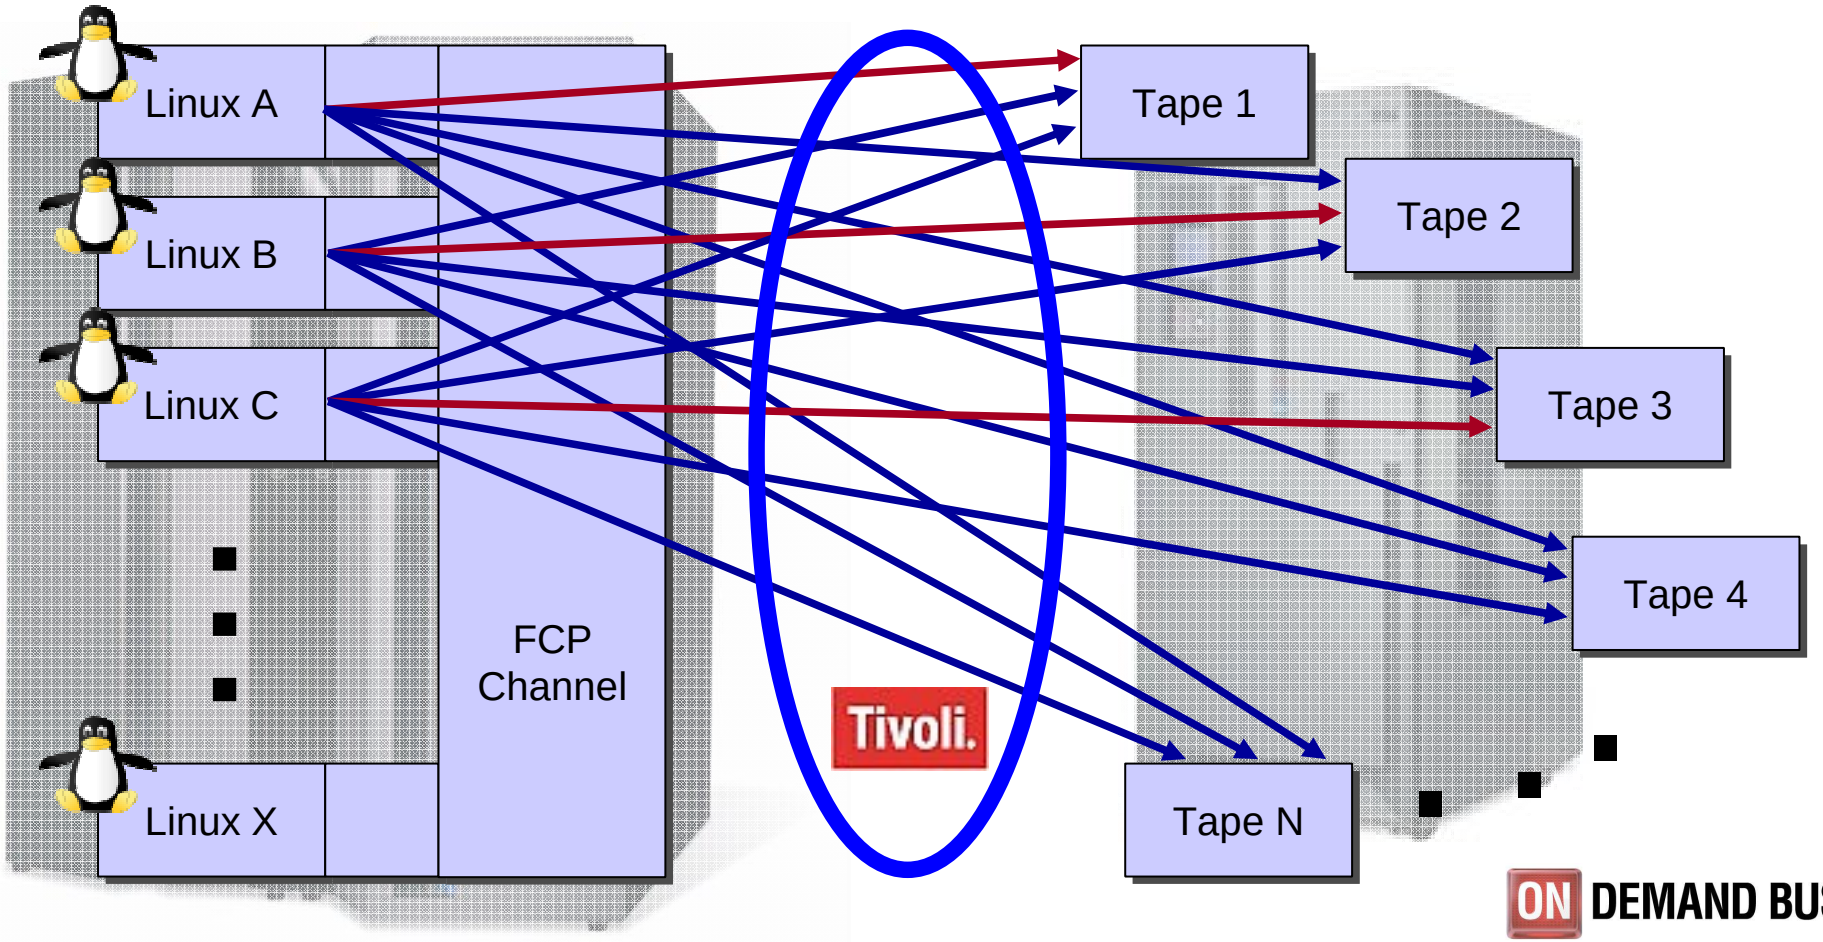

New Possibilities

'Exclusive LUN' policy eliminated

- Shared SAN file systems

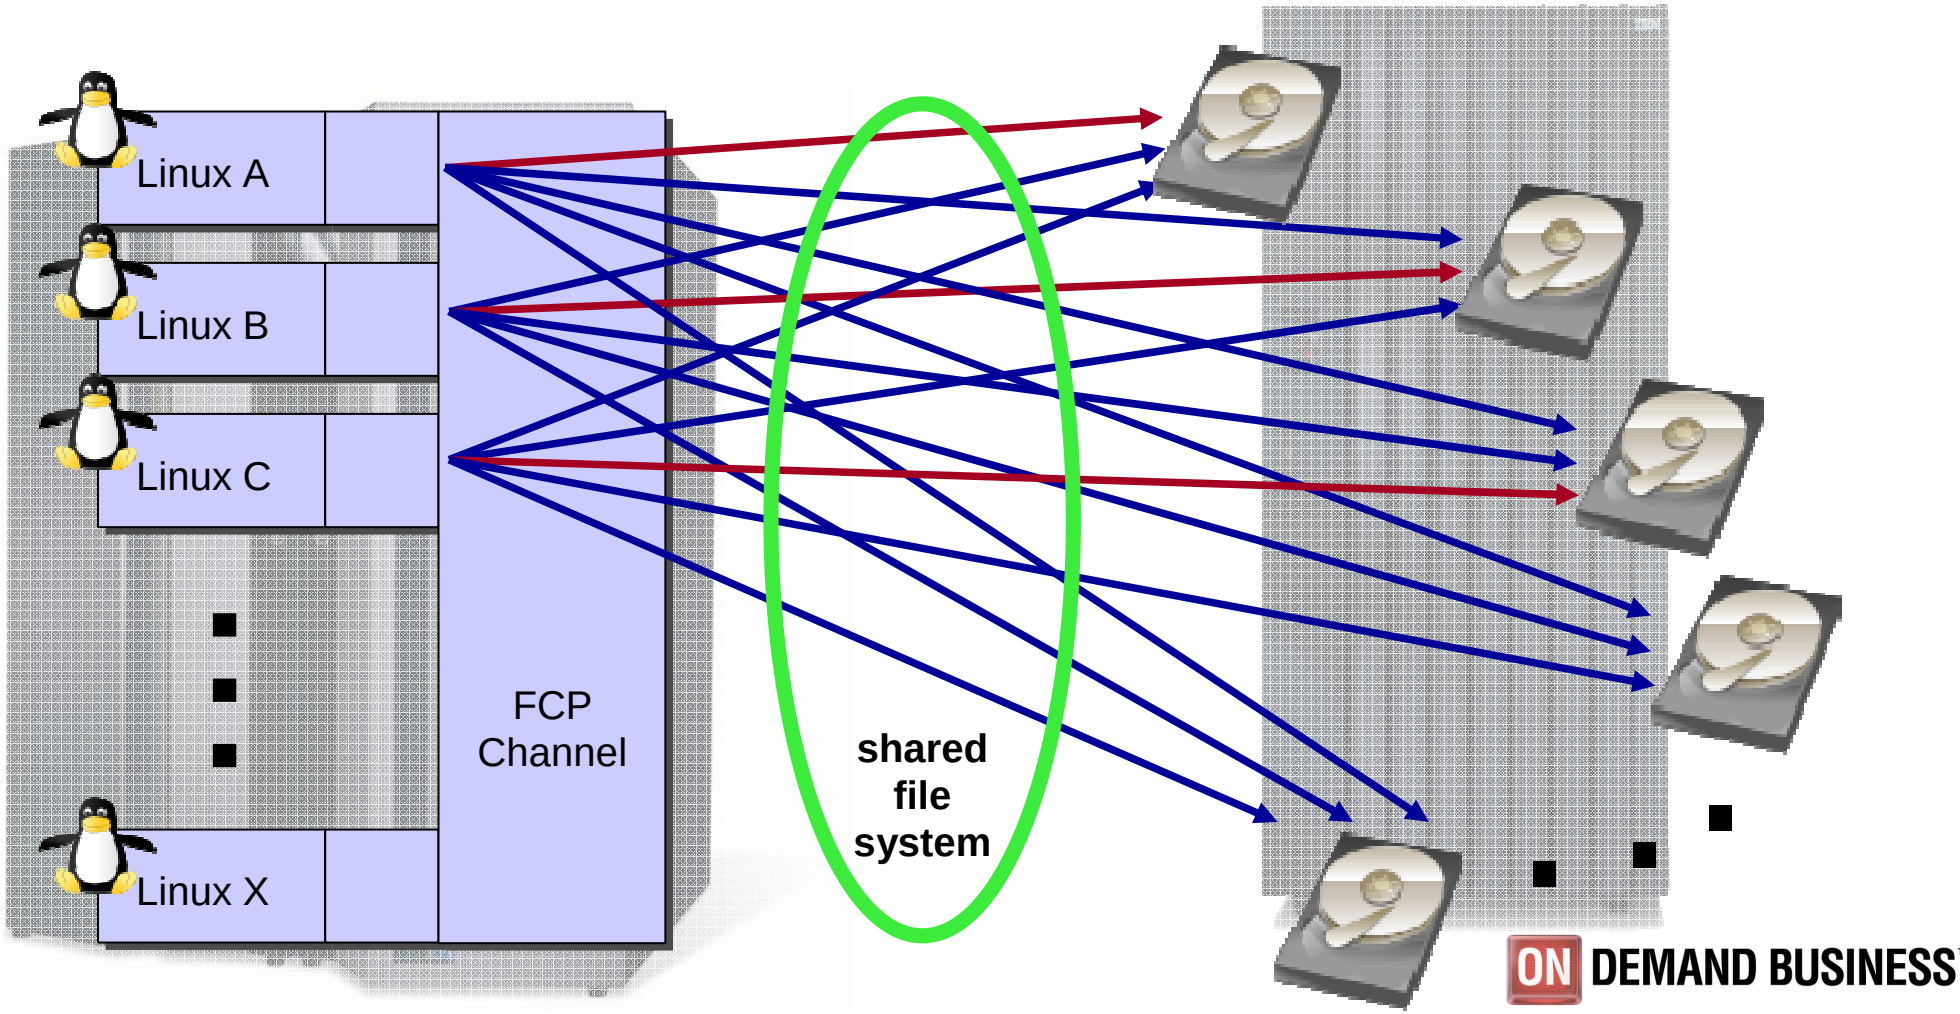

New Possibilities

Access Control Done Right

- LUN Masking works for virtual servers!

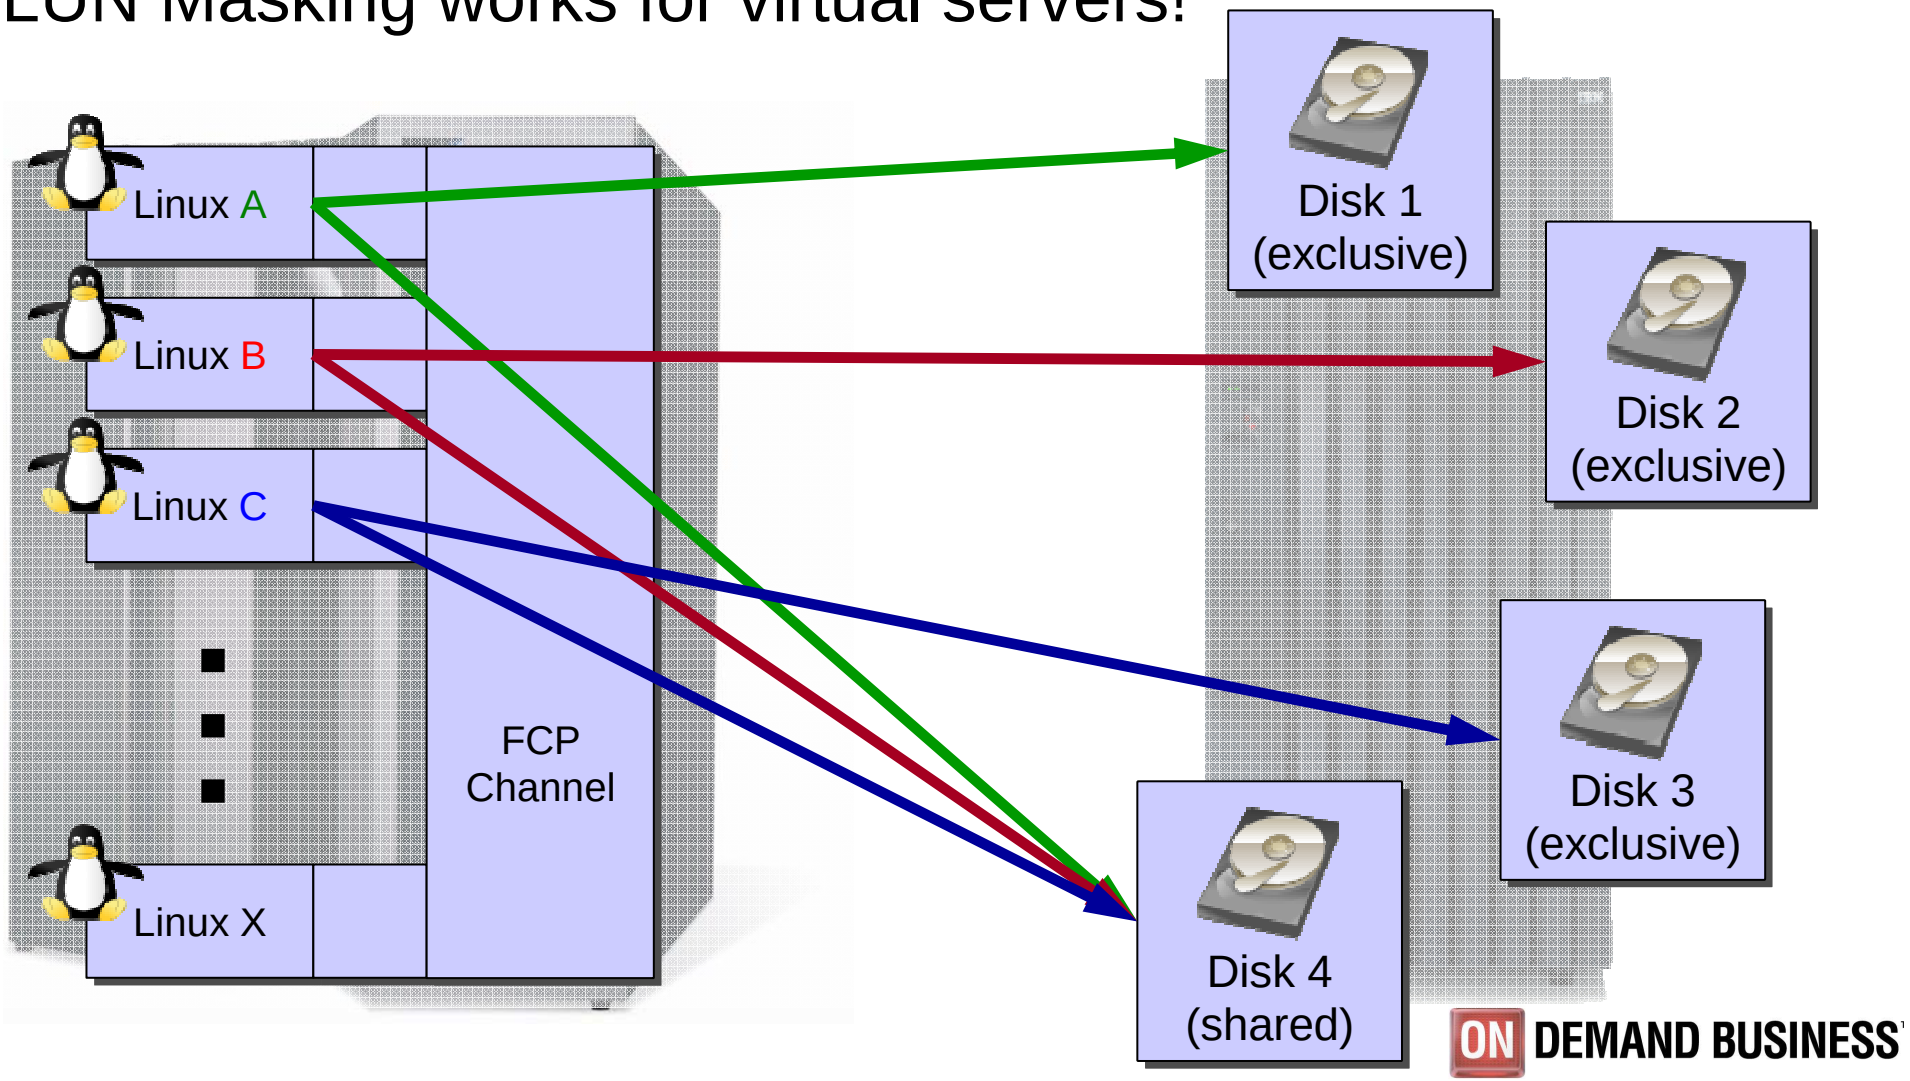

New Possibilities

Access Control Done Right

- Zoning works as well!

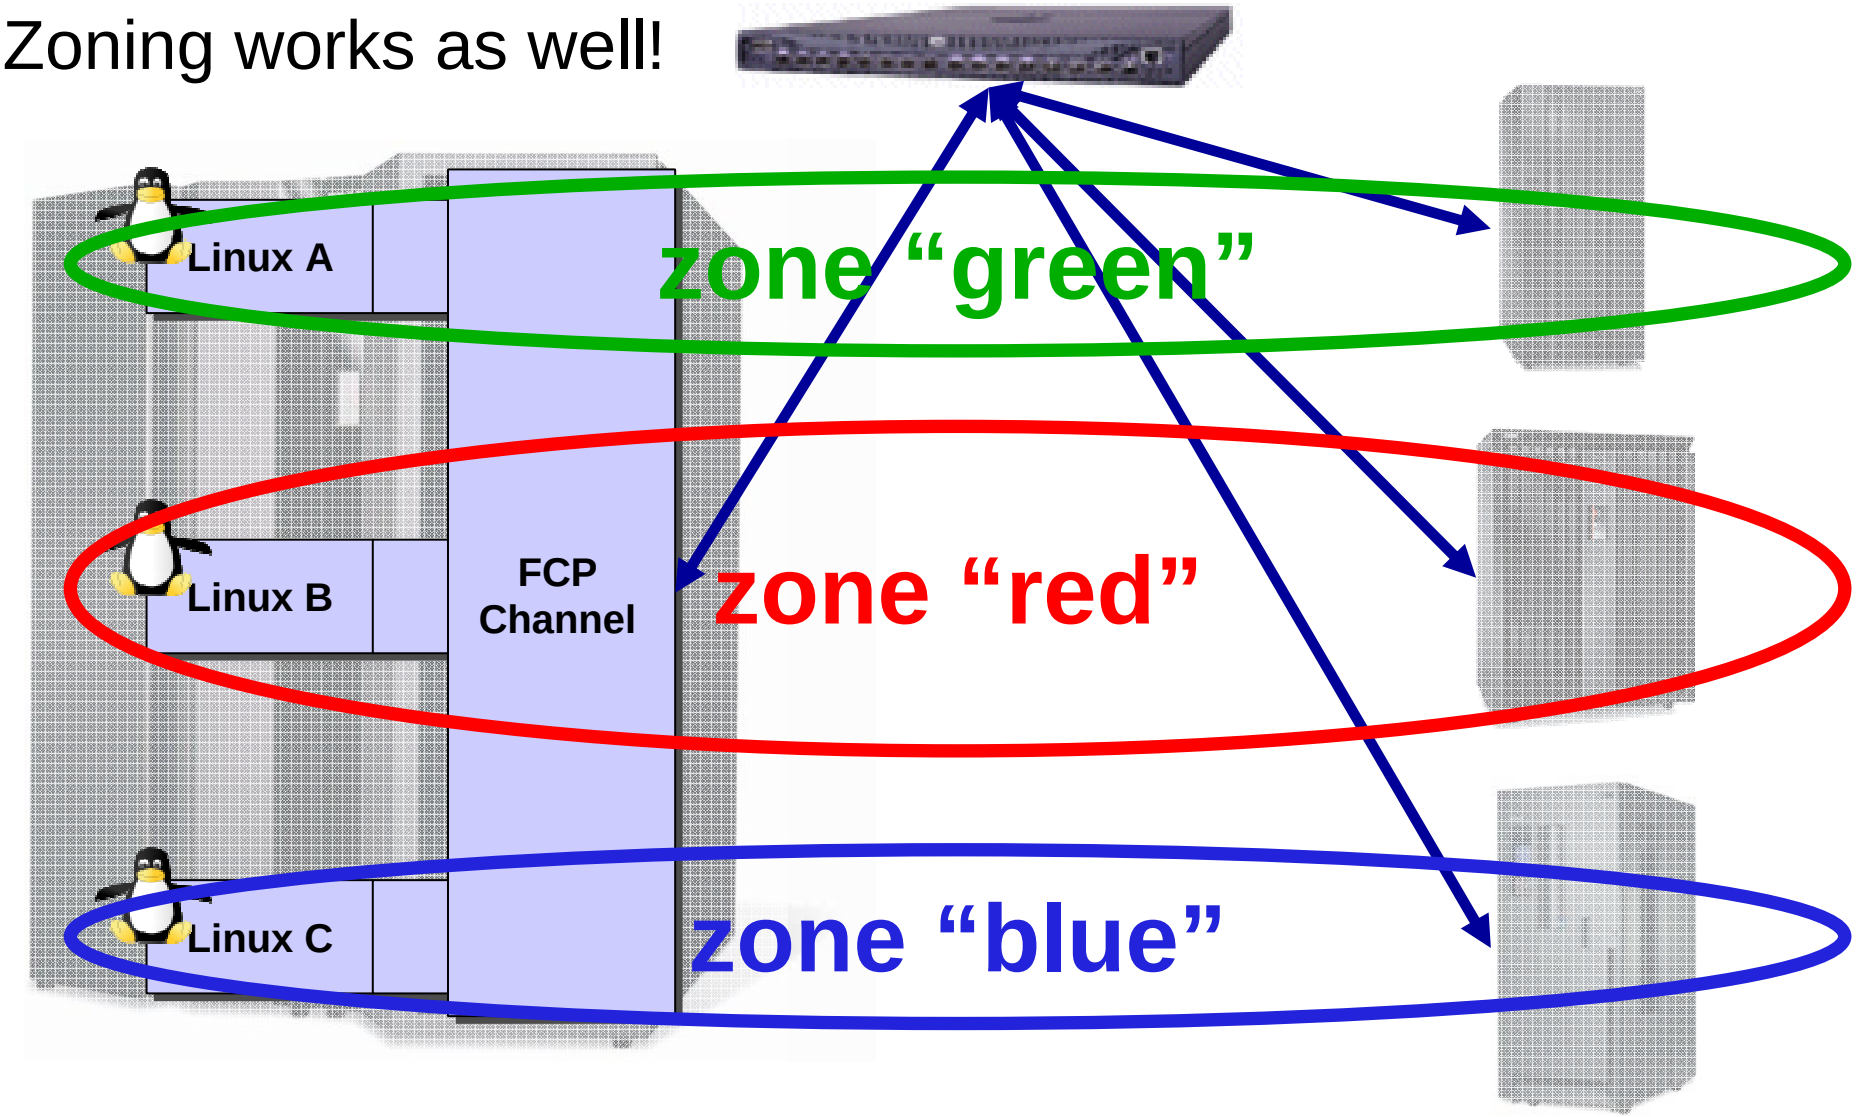

New Possibilities

Access Control Feature

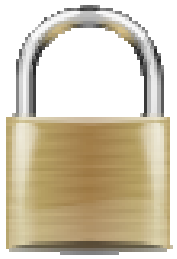

- NPIV deprecates the FCP Channel Access Control feature
- Access Control feature still available with System z9, (could be used for subchannels operating in backward-compatible non-NPIV mode)
- ACT rules – if defined – are not applied to FCP subchannels in NPIV mode

Requirements


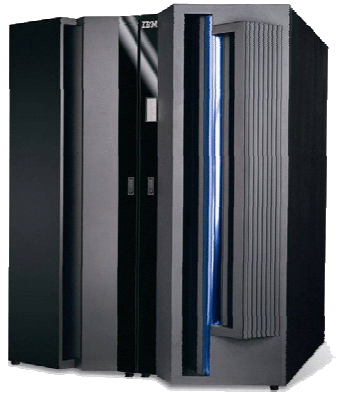

- NPIV is available on System z9 servers.
  - FICON Express 2 adapter running with MCL003 on EC J99658
- z/VM
  - z/VM 5.2
  - z/VM 5.1 with the PTF for APAR VM63744
- Linux Distribution
  - Currently SLES9 SP3 and SLES10 (LPAR mode or z/VM)
- NPIV-Capable Switch
  - only required for switch adjacent to z9
  - Mostly firmware upgrades possible (e.g. McData, Brocade)

Linux Support

Almost Dispensable

- NPIV is more or less transparent for operating systems
  - Linux uses the new virtual N-Port in the same way as it has used non-virtual N-Ports
- But: some new error codes/messages defined for NPIV-type subchannels
  (e.g. if you run out of D_IDs)
- Linux code shipped with SLES9 SP3 and SLES10

see also	Introducing N_Port Identifier Virtualization for IBM System z9 http://www.redbooks.ibm.com/redpapers/pdfs/redp4125.pdf
see also	Linux on zSeries http://www.ibm.com/developerworks/linux/linux390/index.html

Getting Started

“Floods” of WWPNs

- Many new WWPNs to be used by zoning and LUN masking/mapping functions
- Can be exported from SE through FTP

Getting Started

1) Pre-plan SAN with NPIV support
   - see practical limits of components
2) Define FCP subchannels in IOCDS
   - prior to IML, or
   - using dynamic I/O (HCD/HCM)
3) Perform IML, if needed
   - WWPNs for new subchannels get assigned
4) Query WWPNs using SE/HMC panel
   - needed for configuration of SAN functions
   - export function through FTP available

NPIV
Step-by-Step

Getting Started

NPIV Step-by-Step (cont.)

5) Configure switch adjacent to z9
   - ensure NPIV is enabled
   - ensure enough virtual N-Ports per port
   - setup zoning for virtual N-Ports
6) Configure target device
   - setup LUN masking/mapping
7) Enable NPIV-mode for CHPID in LPAR
   - CHPID must be temporarily toggled off
8) Start using FCP subchannel in Linux
   - check for NPIV related error messages

Getting Started

Implementation Limits

- Switches will presumably allow for 1024 or up to 8192 WWPNs in a SAN
- Storage device impose limits as well
- Theoretically up to 255 subchannels per channel connected at the same time
- ≤ 510 active target port connections for all subchannels of a channel
- Each System z9 provides a total of 2 million WWPNs for virtual HBAs

see also	Support of Fiber Channel Protocol for SCSI - FCP channels http://www-03.ibm.com/systems/z/connectivity/fcp.html

Hints and Tips

- Do not use more than 32 subchannels per physical channel in NPIV mode.
- Zone each NPIV WWPN individually. This can reduce fabric traffic.
- Consider using multipathing (performance and availability).
- Enable NPIV on the SAN switch before enabling it on the System z9 server.
- Be aware that each login from a NPIV-mode subchannel into a storage subsystem counts as a separate host login. There are limits at storage side.
- Switches typically limit the number of supported N_Port IDs.
- Some switches limit the number of N_Port IDs that can be assigned to a physical port.
- FCP microcode MCL003 on EC J99658 requires a special activation procedure. All FCP PCHIDs should be configured off *before* activating the MCL.

Look'n Feel

NPIV Is On

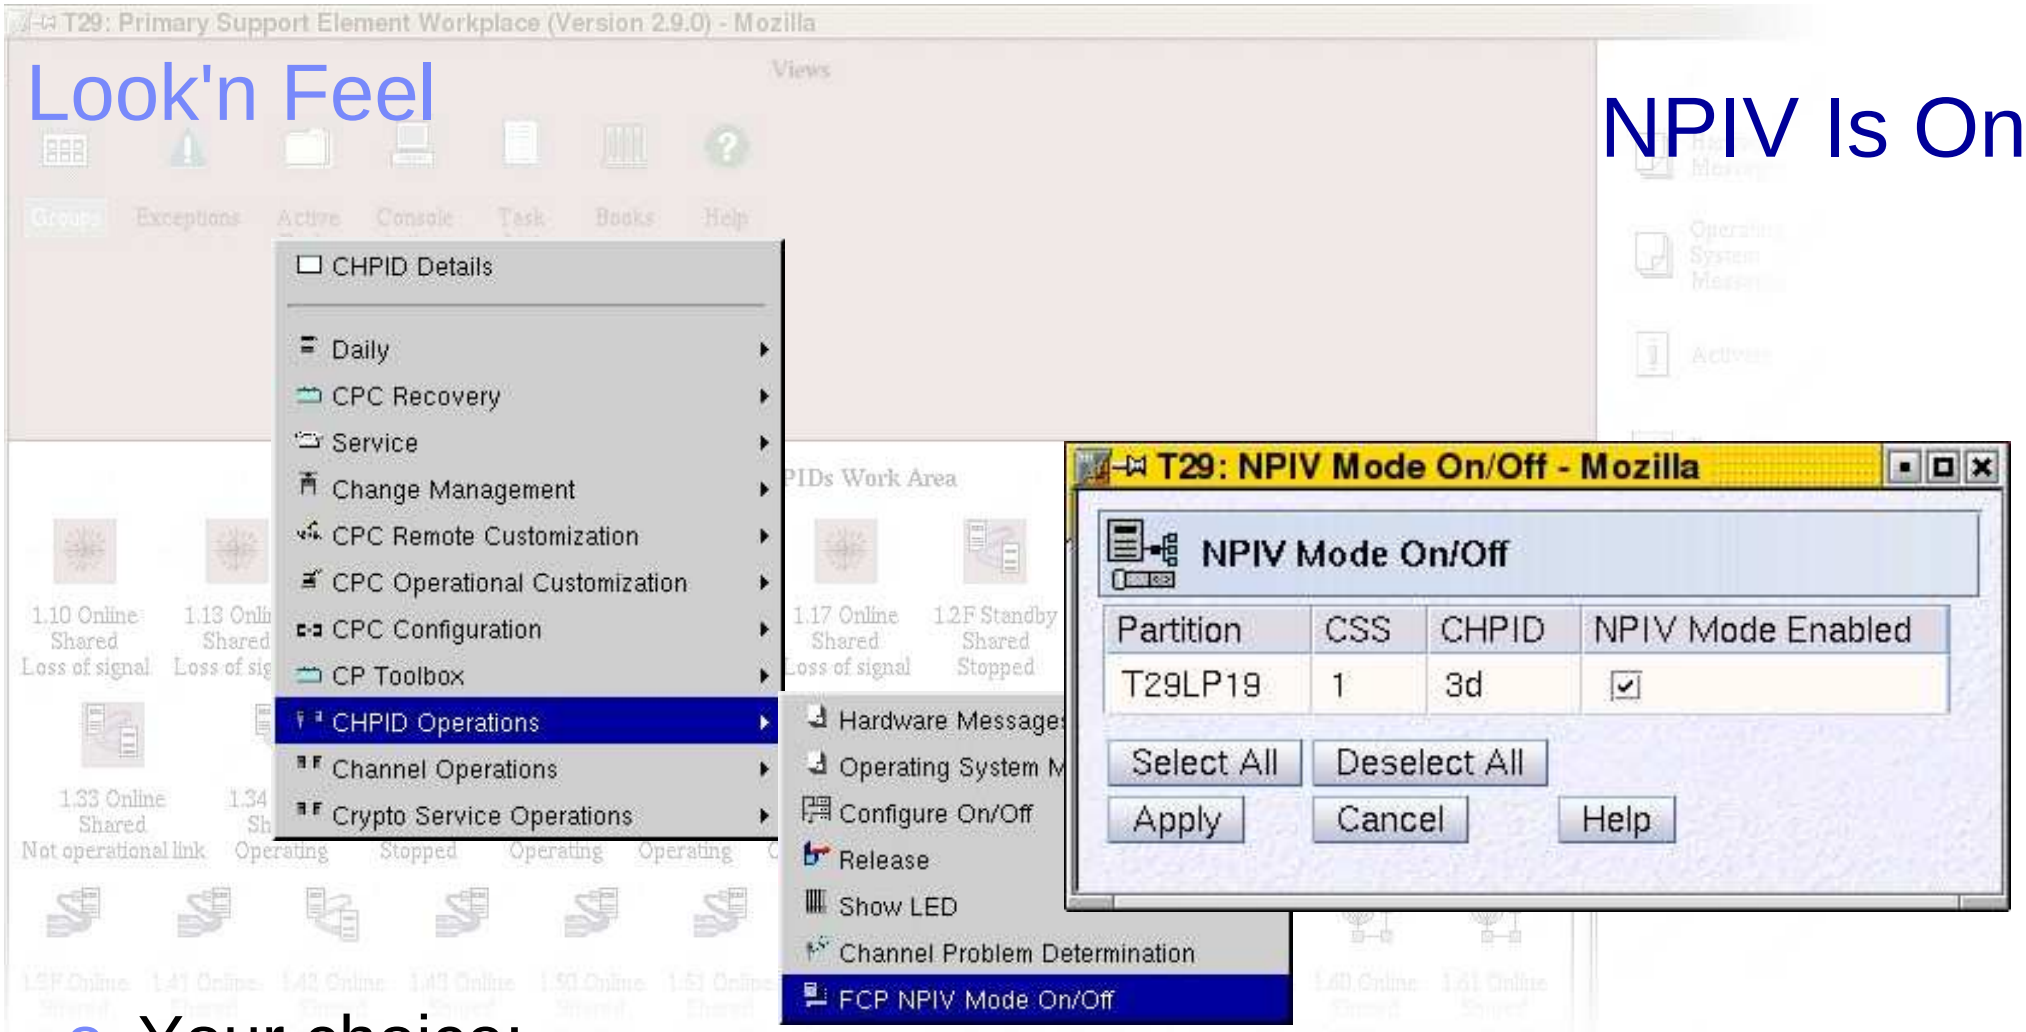

- Your choice:

  toggle NPIV on with a per channel-and-LPAR granularity

  (CHPID needs to be toggled off temporarily)

Look'n Feel

You won't run out of WWPNs, but if ...

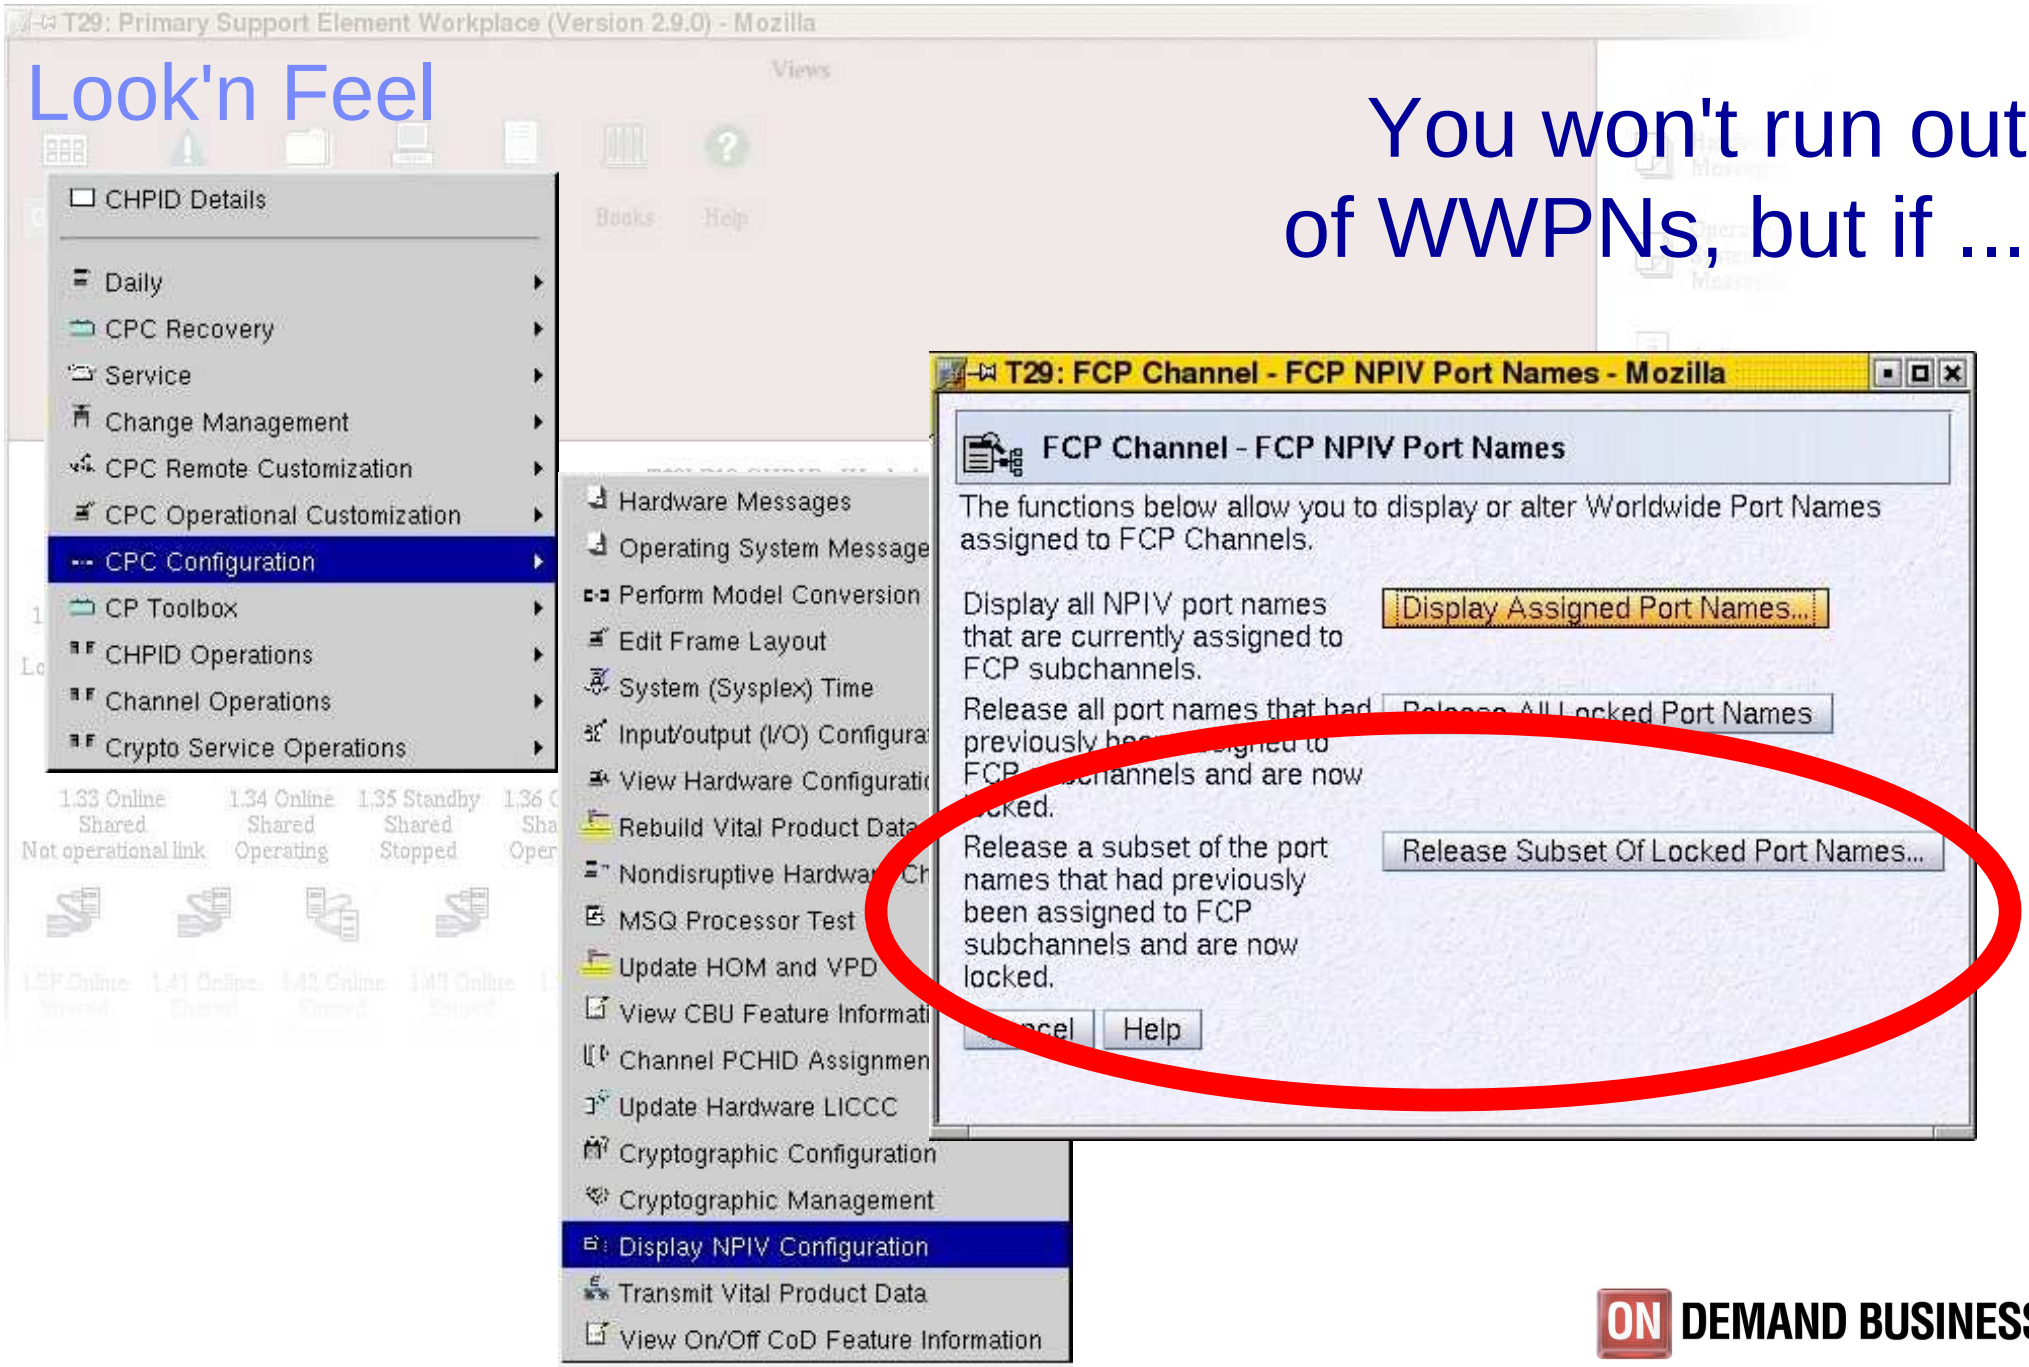

Look'n Feel

List Of Virtual N-Ports

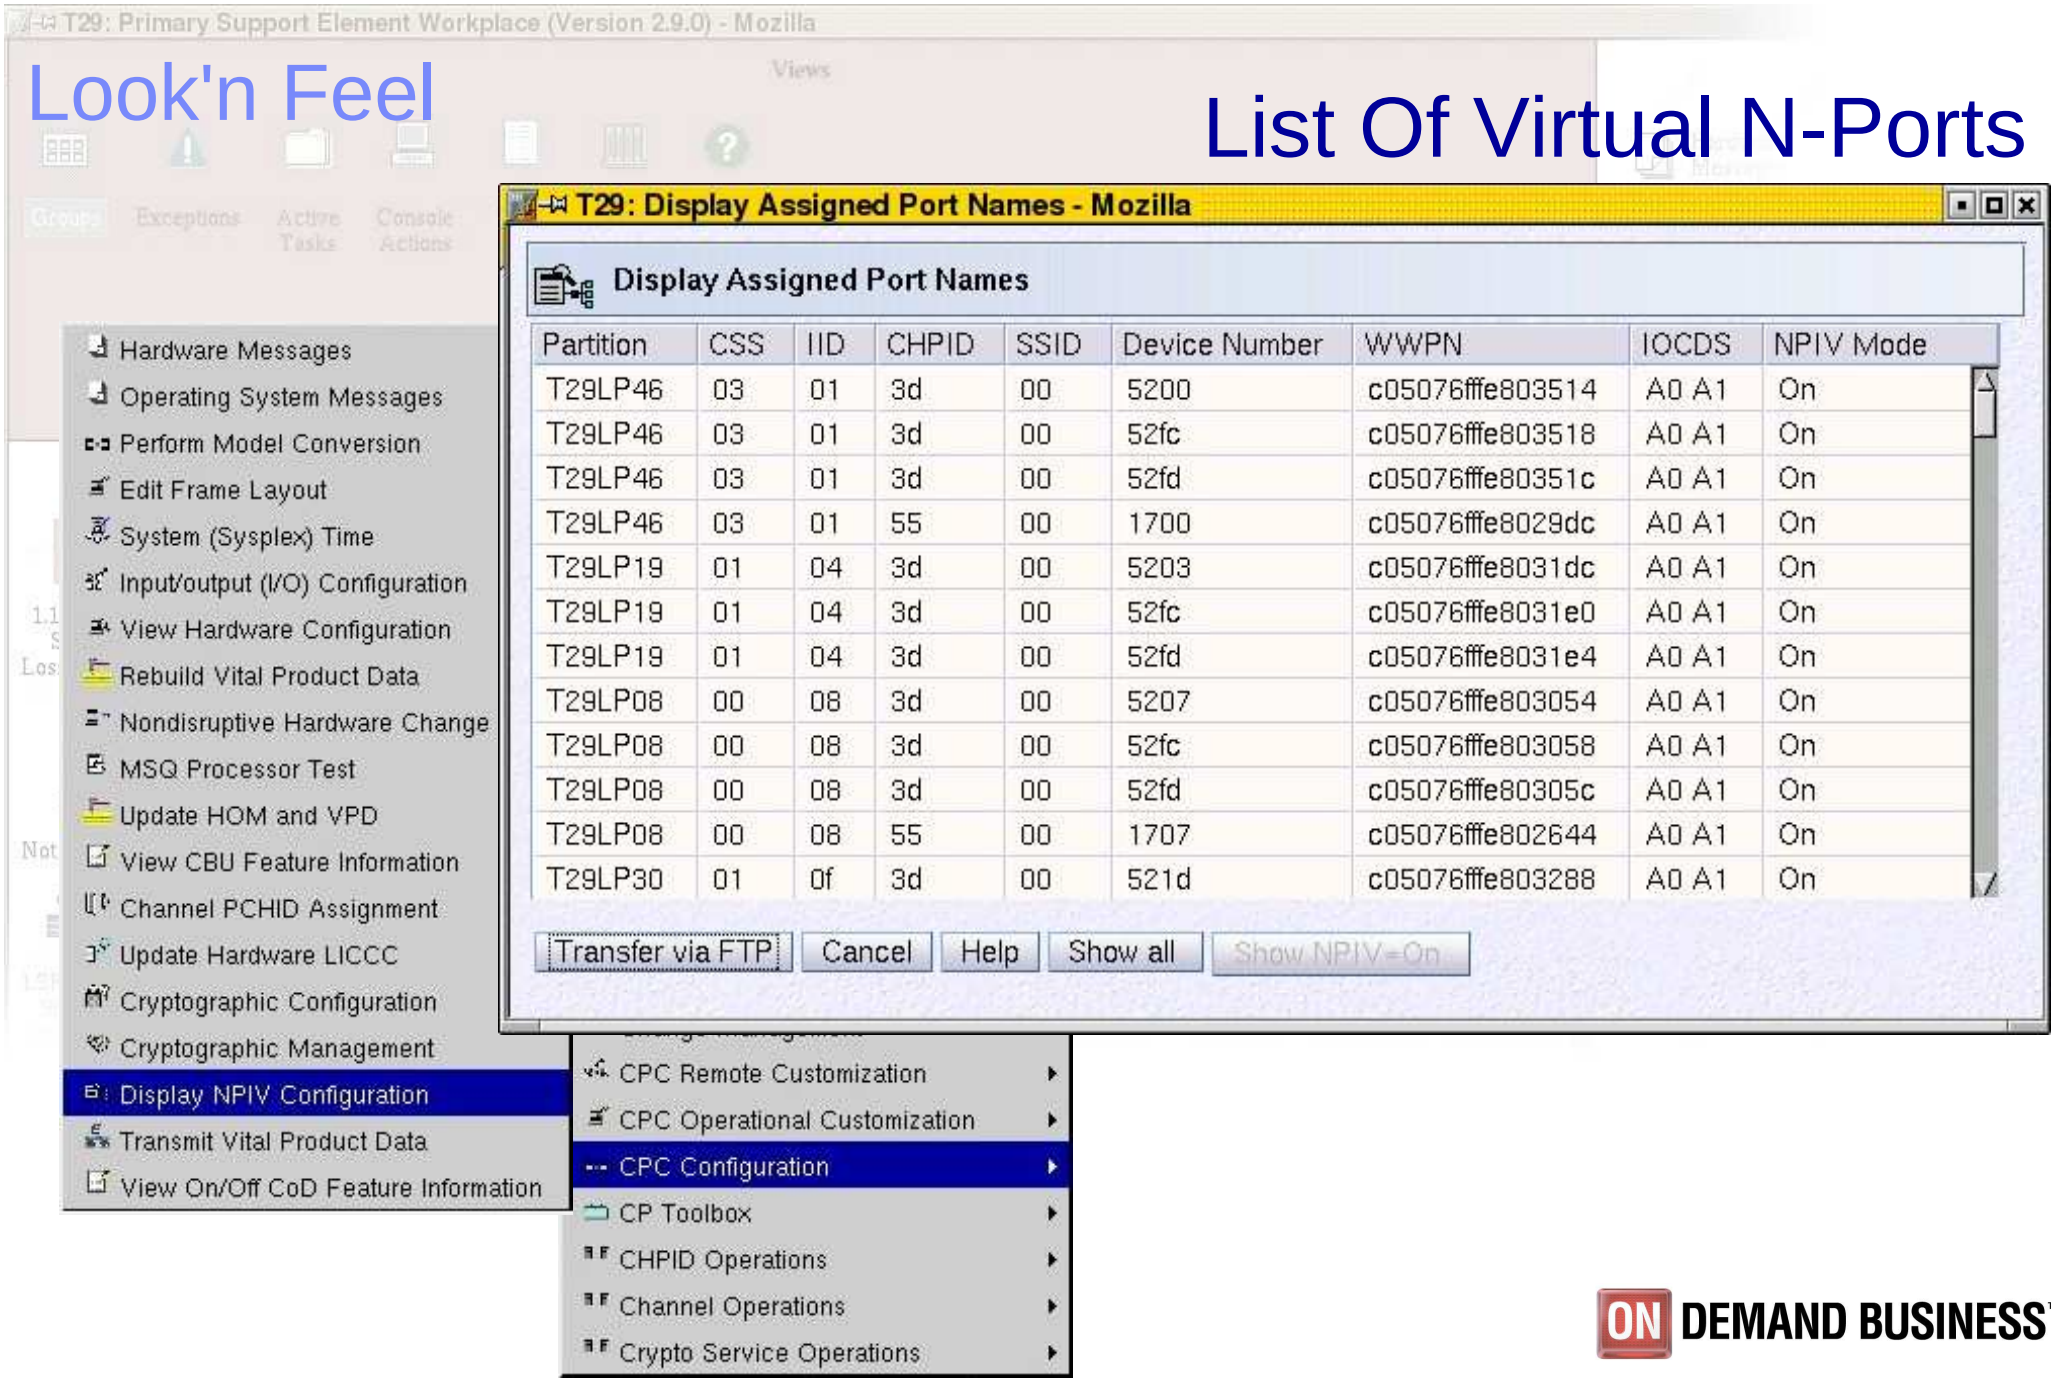

Look'n Feel

Linux Proudly Presents: Its Own SAN Identity

virtual N-Port tied to this FCP subchannel

in non-NPIV mode:

same WWPN

in NPIV mode:

WWPNs differ

physical N-Port tied to FCP channel (CHPID)

```
# pwd
/sys/class/fc_host/host0
```

```
# cat node_name
0x5005076400cd6aad

# cat port_name
0x500507640 1008fa8

# cat permanent_port_name
0x500507640 1008fa8
```

(WWNN - Host)

(WWPN – virtual Adapter)

(WWPN – physical Adapter)

```
# cat node_name
0x5005076400cd6aad

# cat port_name
0xc05076fffe8031dc

# cat permanent_port_name
0x500507640 1008fa8
```

SLES10

Keep in Mind

NPIV is NOT about ...

- … reducing SAN maintenance work
  - but it straightens out the way of doing things
- … migrating virtual servers more easily
  - but it enables virtual servers in FC SANs
- … introducing new FC security schemes
  - but it integrates virtual servers into existing ones
- … solving any SAN issue immediately
  - but it provides a foundation to build on

Summary

„Ready, steady, ... go!“

- NPIV removes the root cause for limited practicability of FCP in virtual server environments
- NPIV allows the use of standard SAN management solutions on IBM System z9
- NPIV gives free rein to sophisticated SAN fantasies

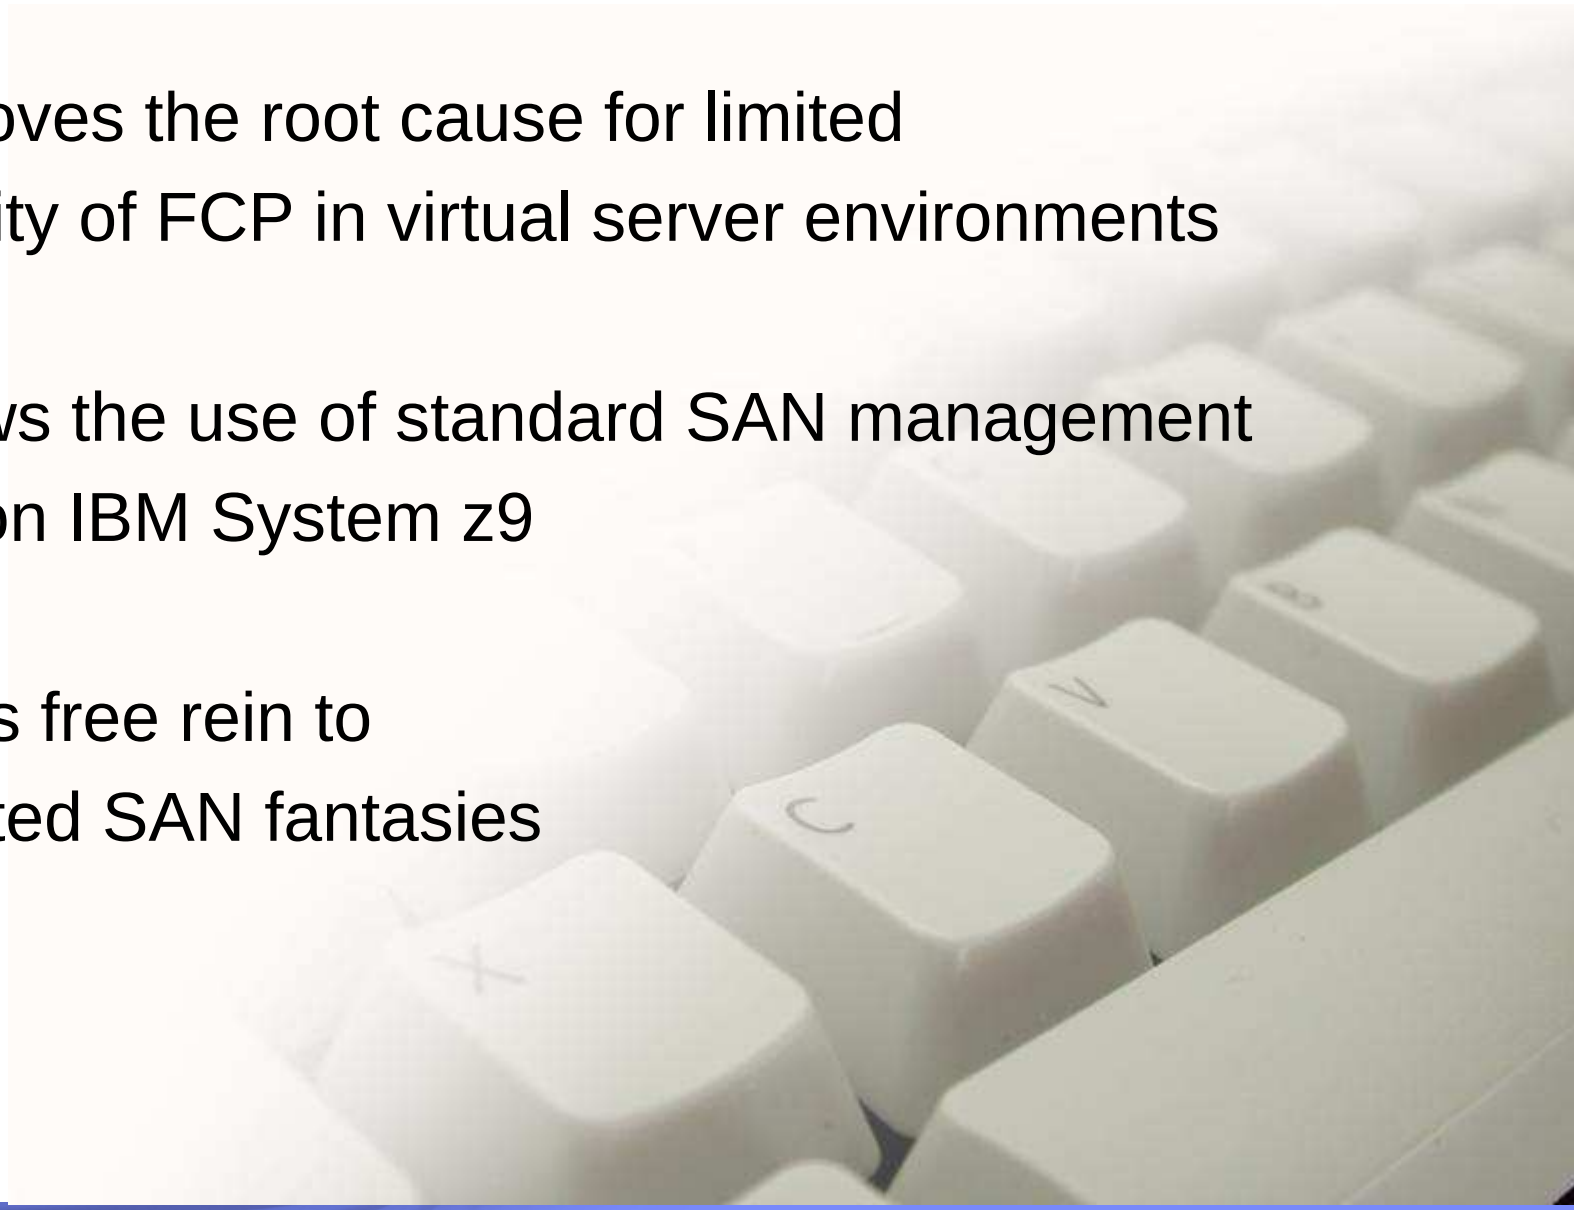

NPIV
FCP Channel Virtualization in a Linux Environment

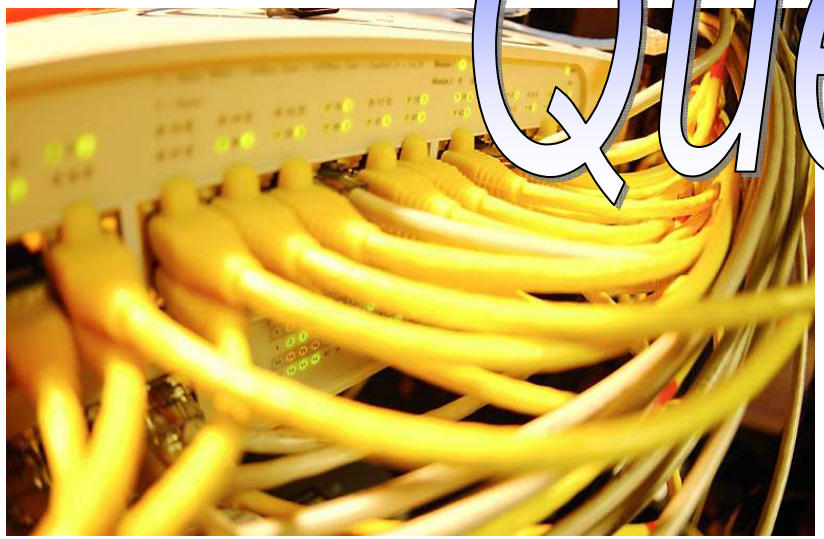

Questions ?

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries:

Enterprise Storage Server, IBM*, IBM logo*, System z9*, IBM eServer, z/VM, zSeries

*Registered trademarks of IBM Corporation

Linux is a registered trademark of Linus Torvalds.

All other products may be trademarks or registered trademarks of their respective companies.

NPIV
FCP Channel Virtualization in a Linux Environment
Volker Sameske

ON DEMAND BUSINESS