IBM GLOBAL SERVICES

# L93

## FCP Channel Virtualization in a Linux Environment

Martin Peschke, IBM Lab Boeblingen


**September 19 - 23, 2005**

**San Francisco, CA**

# FCP Channel Virtualization

## Implementing Linux with *FCP for zSeries*

zSeries Expo, San Francisco, September 19-23 2005
Martin Peschke, Linux on zSeries Development
mpeschke@de.ibm.com

# Trademarks



# Acknowledgments

**Thanks for suggestions and input to Gerhard Banzhaf, Stefan Müller, Maxim Shchetynin**

**Implementing Linux with FCP for zSeries - FCP Channel Virtualization**
**Martin Peschke**

# FCP Channel Virtualization

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

# What is NPIV?

## Unique SAN Identities!

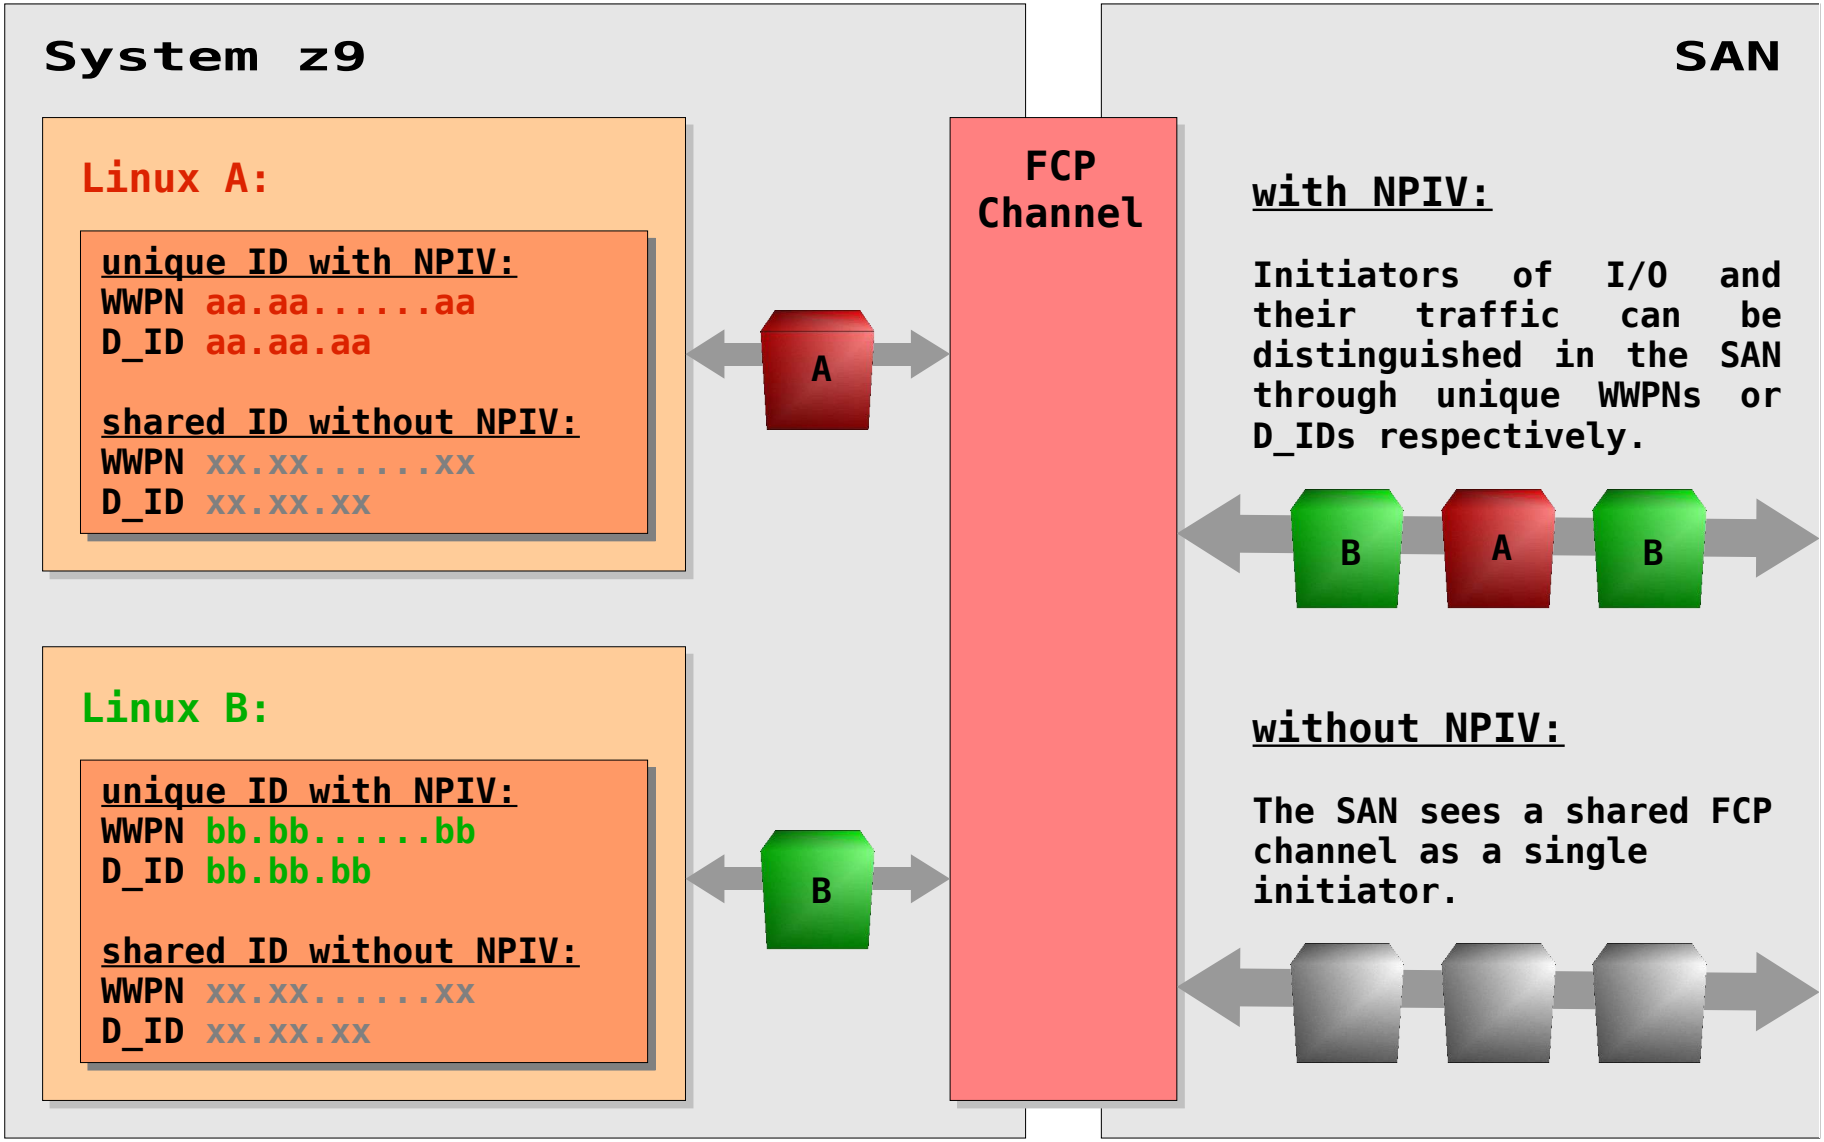

# What is NPIV?

## Think Virtual Adapters

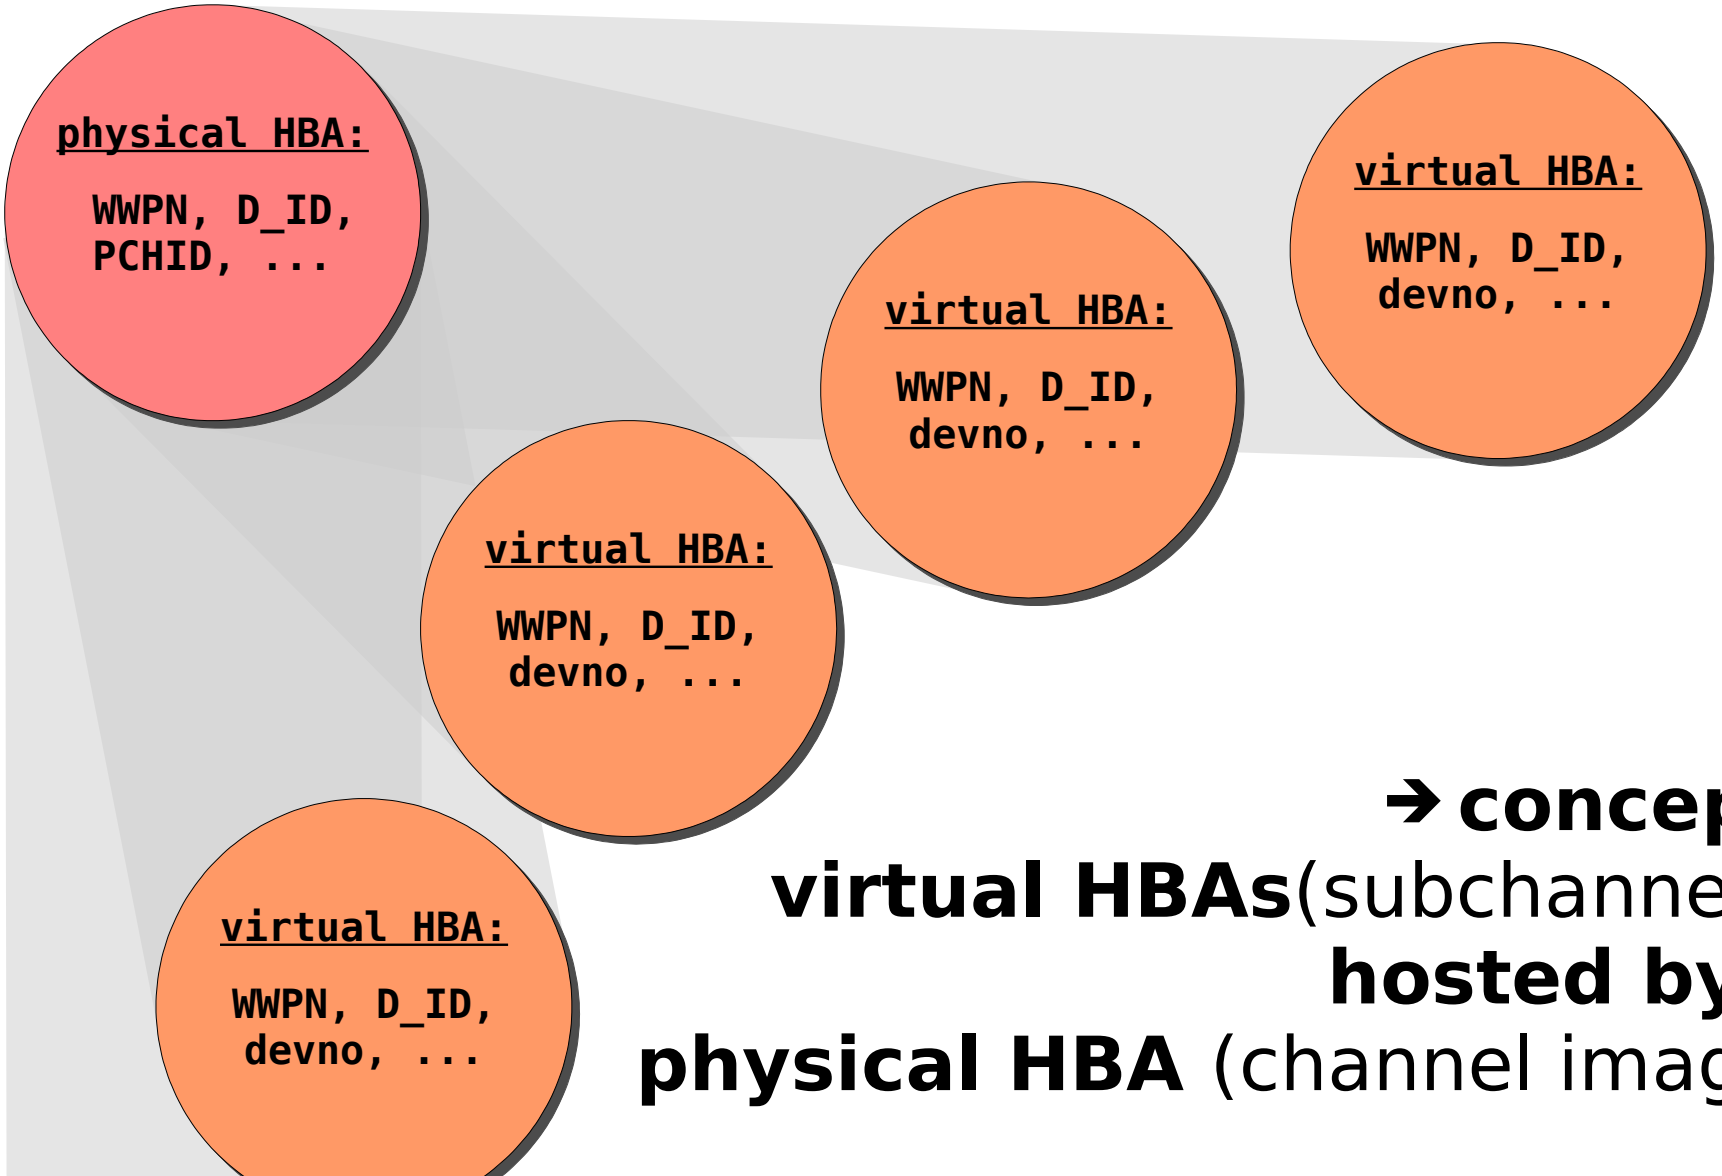


**→ concept:**
**virtual HBAs**(subchannels)
**hosted by a**
**physical HBA** (channel image)

# 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 *Fibre Channel Framing and Signaling Interface* http://www.t11.org

# What is NPIV?

## System z9 Implementation

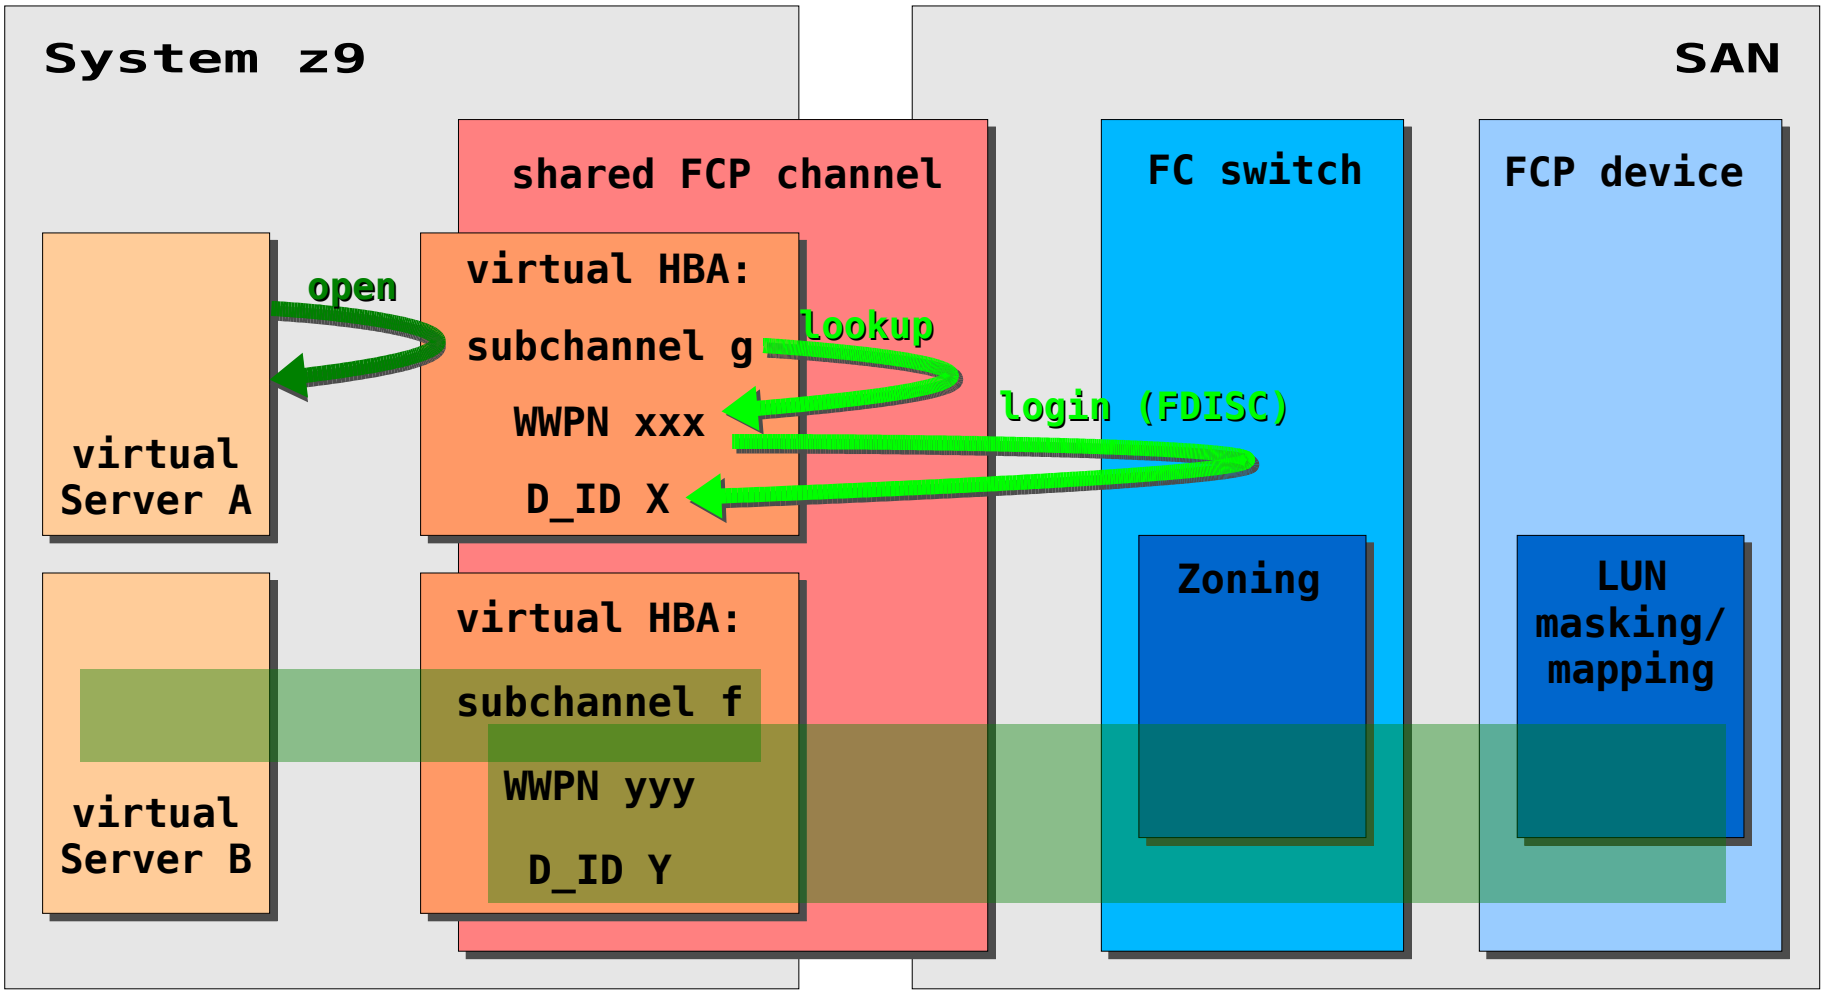


→ connecting a virtual HBA to the SAN

# What is NPIV?

## System z9 Implementation

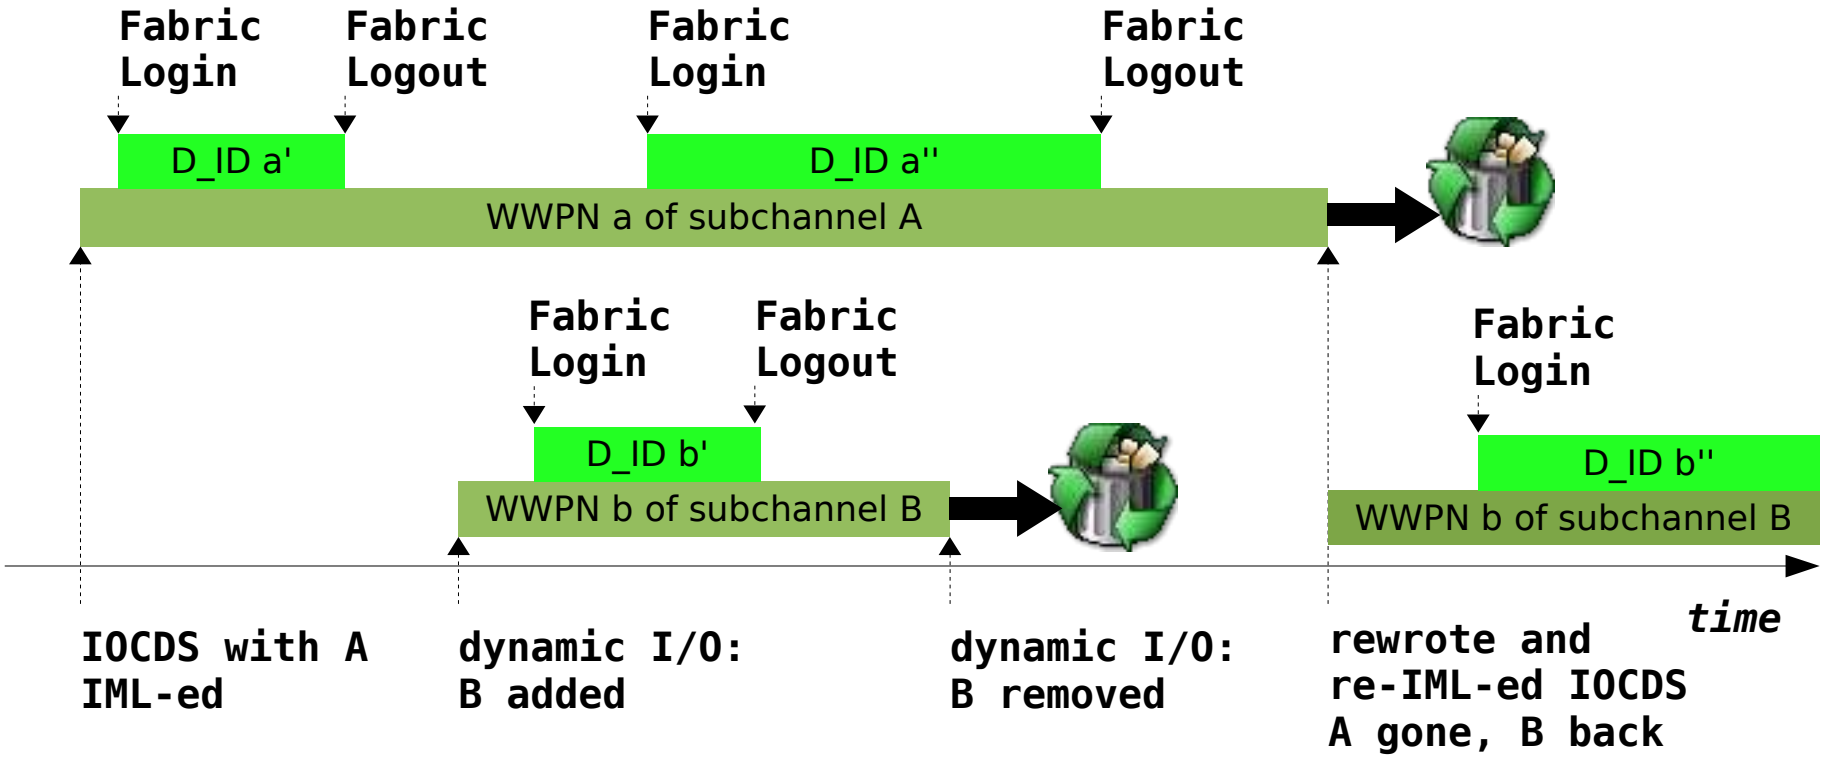


- ➔ **lifetime of virtual WWPN= lifetime of subchannel definition**
- ➔ **D_ID lifetime = fabric connection lifetime**

# What is NPIV?

## System z9 Implementation

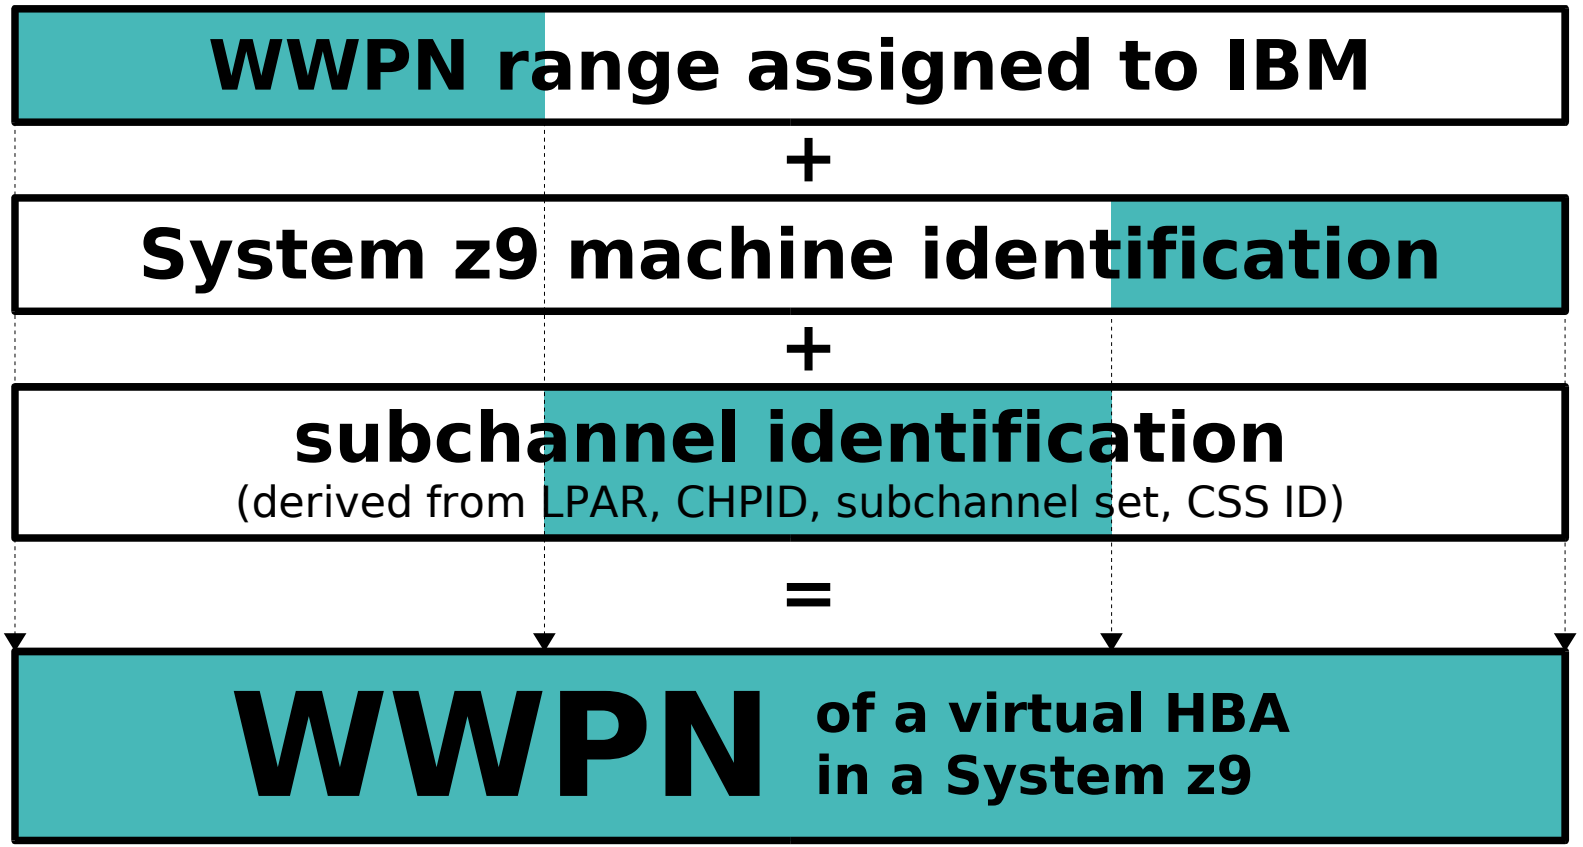


→ **well-defined “ingredients” guarantee uniqueness and permanence of WWPNs**

# New Possibilities

## Sharing Unlimited

➔ **'Exclusive LUN' policy abolished by NPIV: many-to-many tape backup solutions**

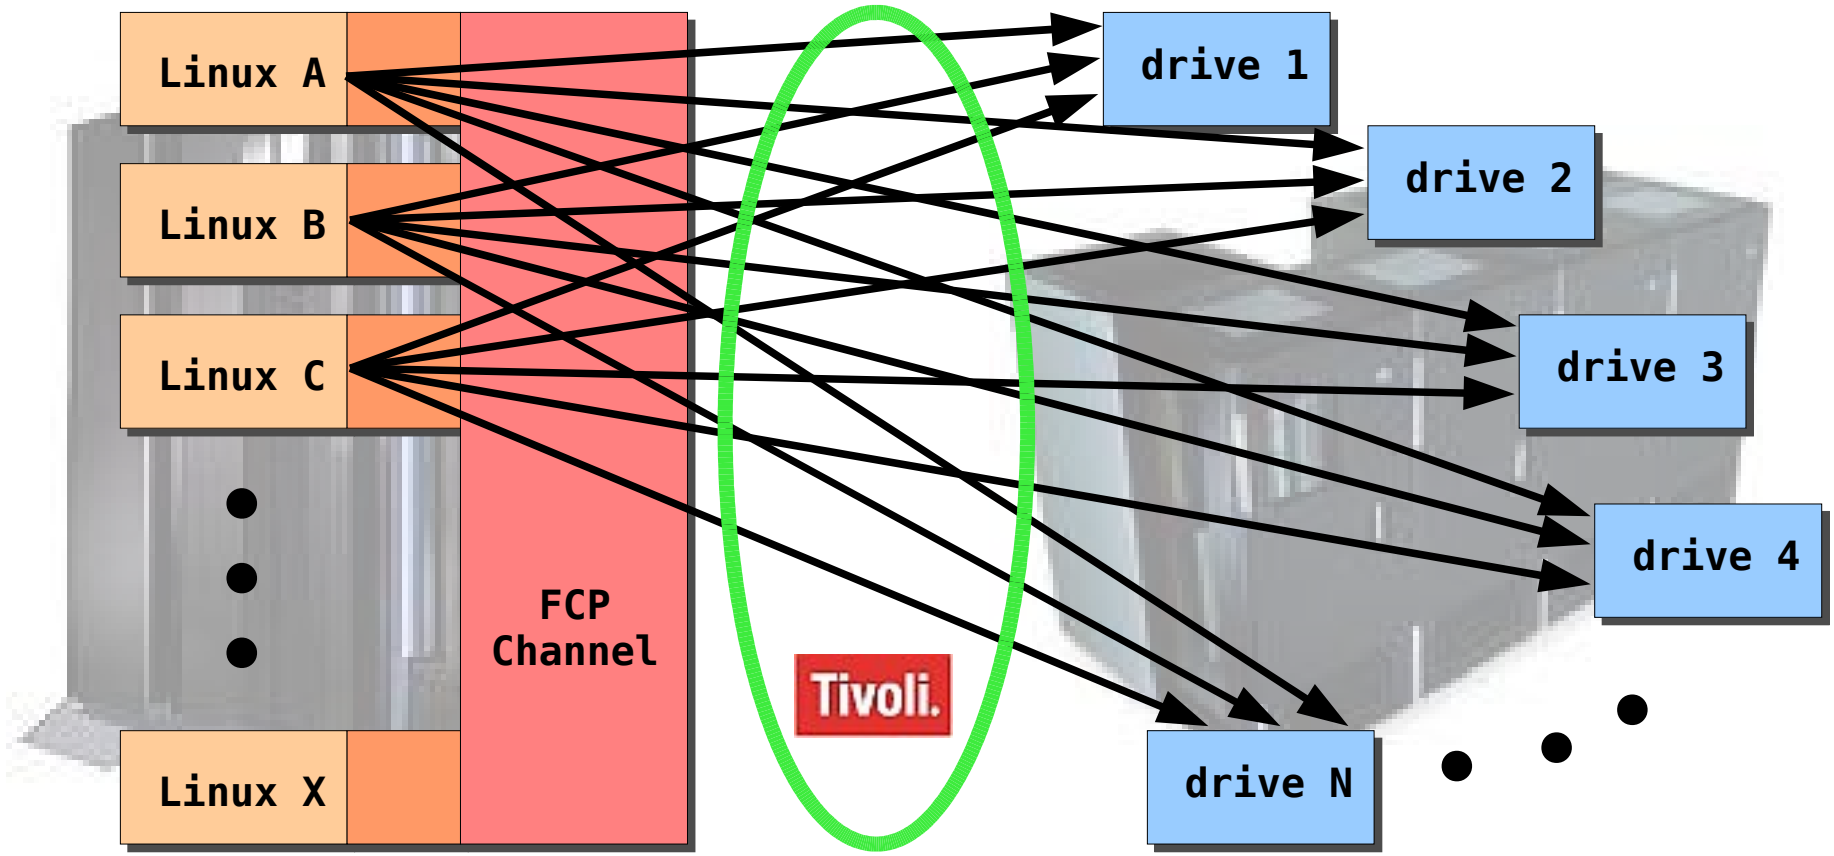

# New Possibilities

## Sharing Unlimited

➔ **'Exclusive LUN' policy abolished by NPIV: imagine shared SAN filesystems**

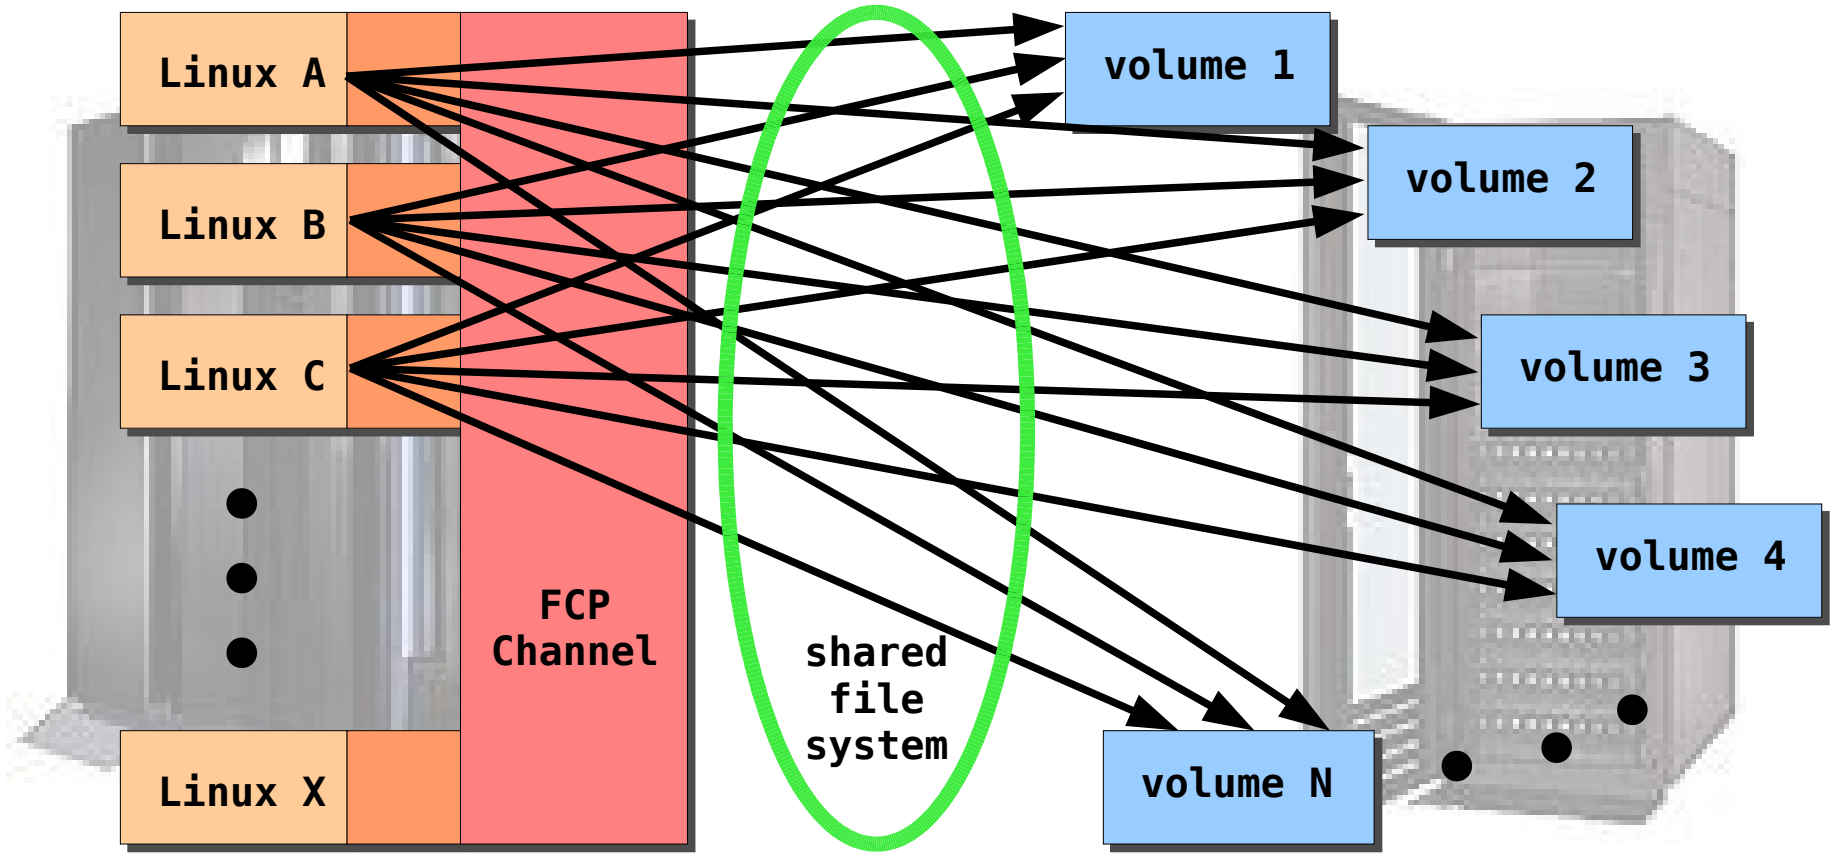

# New Possibilities

## Access Control Done Right

➔ **LUN Masking works for virtual servers!**

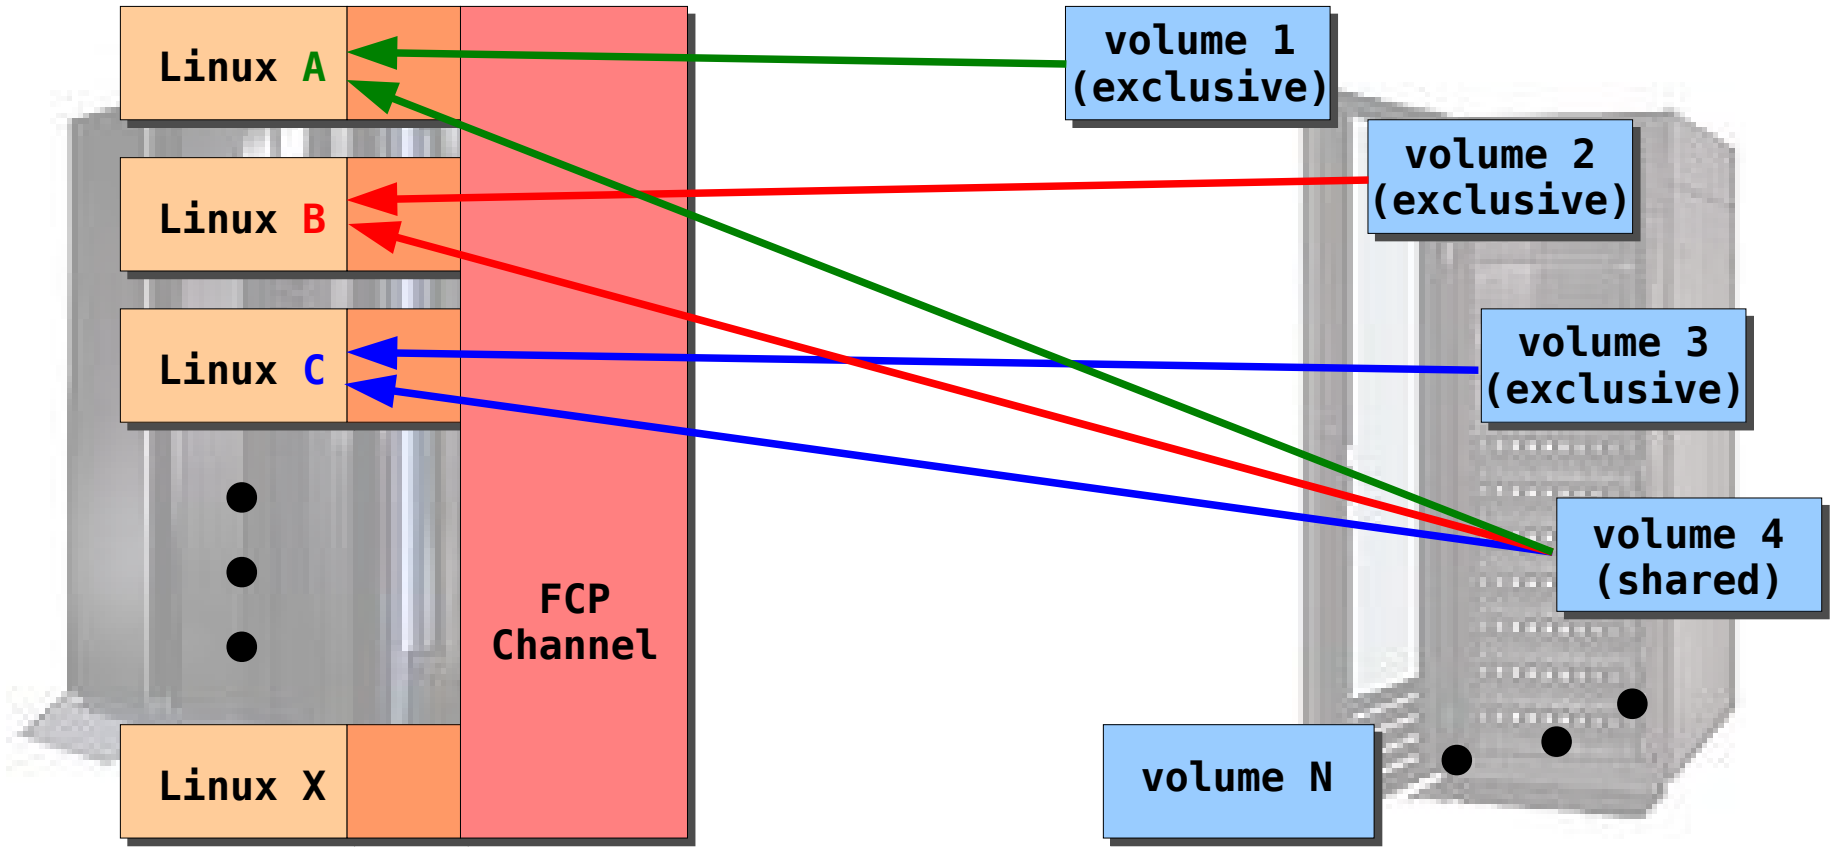

# New Possibilities

## Access Control Done Right

➔ **Zoning works as well!**

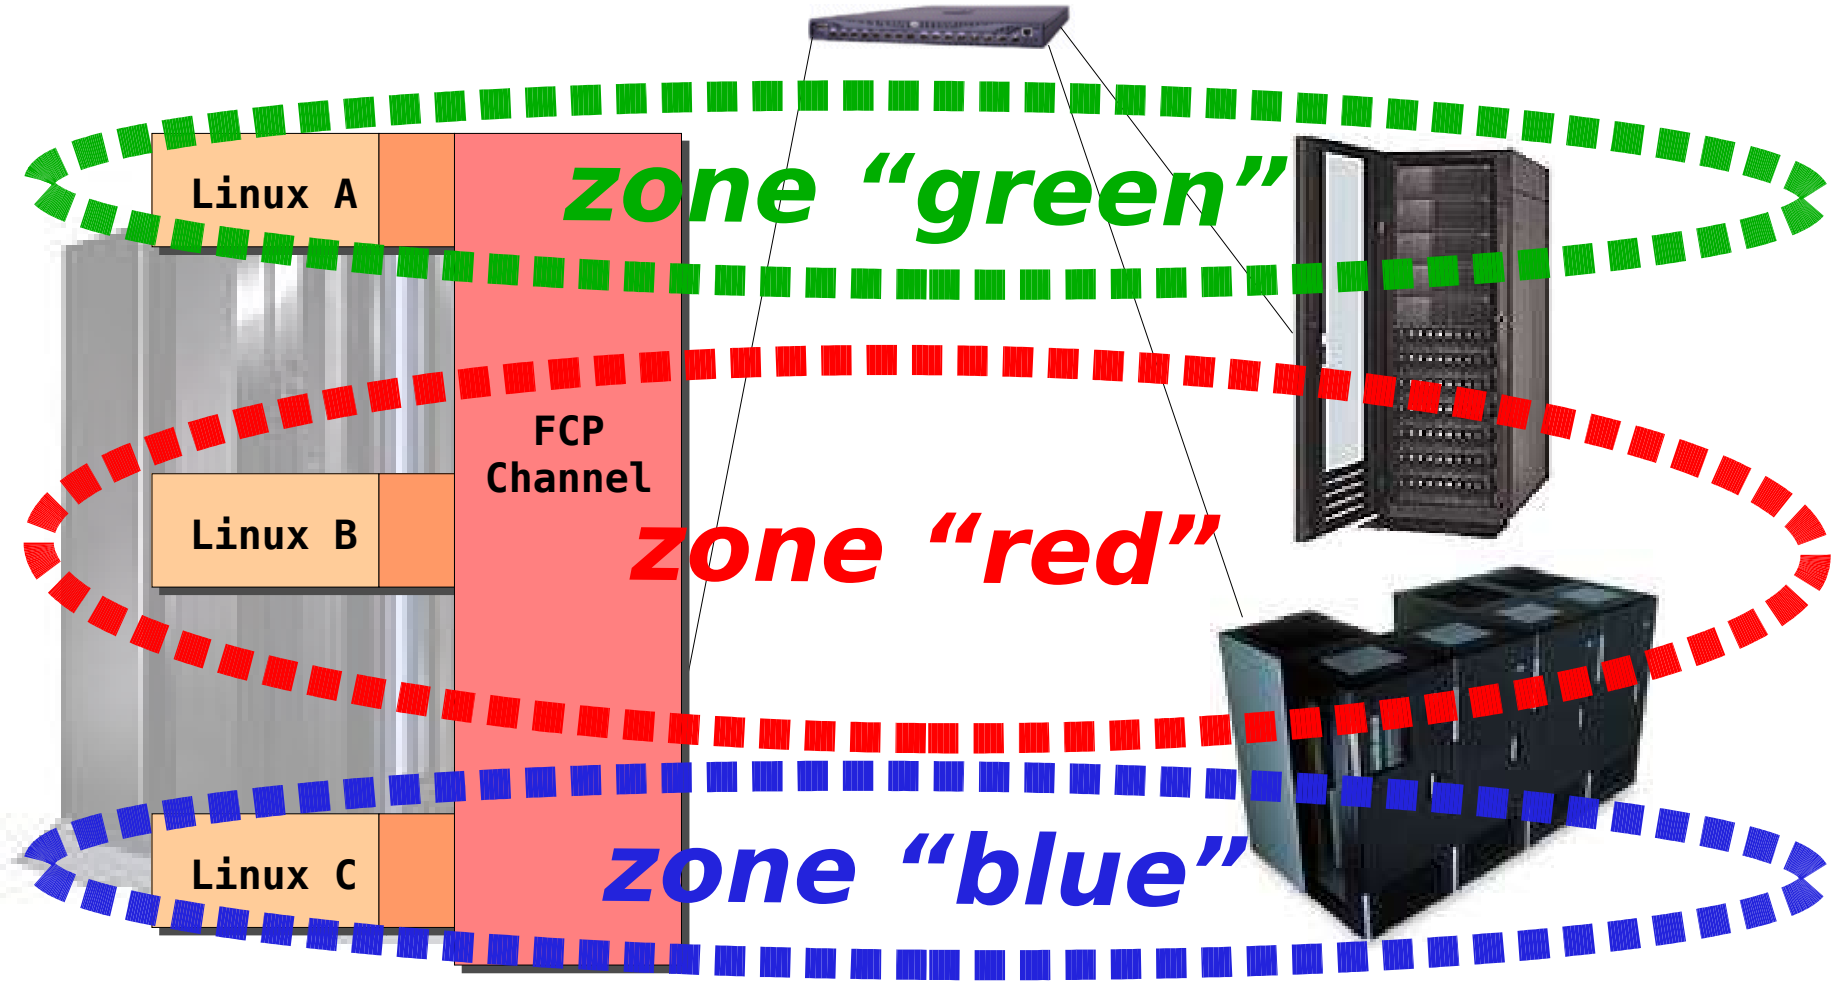

# New Possibilities

## Access Control Done Right

- **NPIV deprecates the FCP Channel Access Control feature**
- **Access Control feature still available with System z9, though**
  (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**

# New Possibilities

## All-Out Manageability

➔ **SCSI storage provisioning for virtual servers created by cloning – why not?**

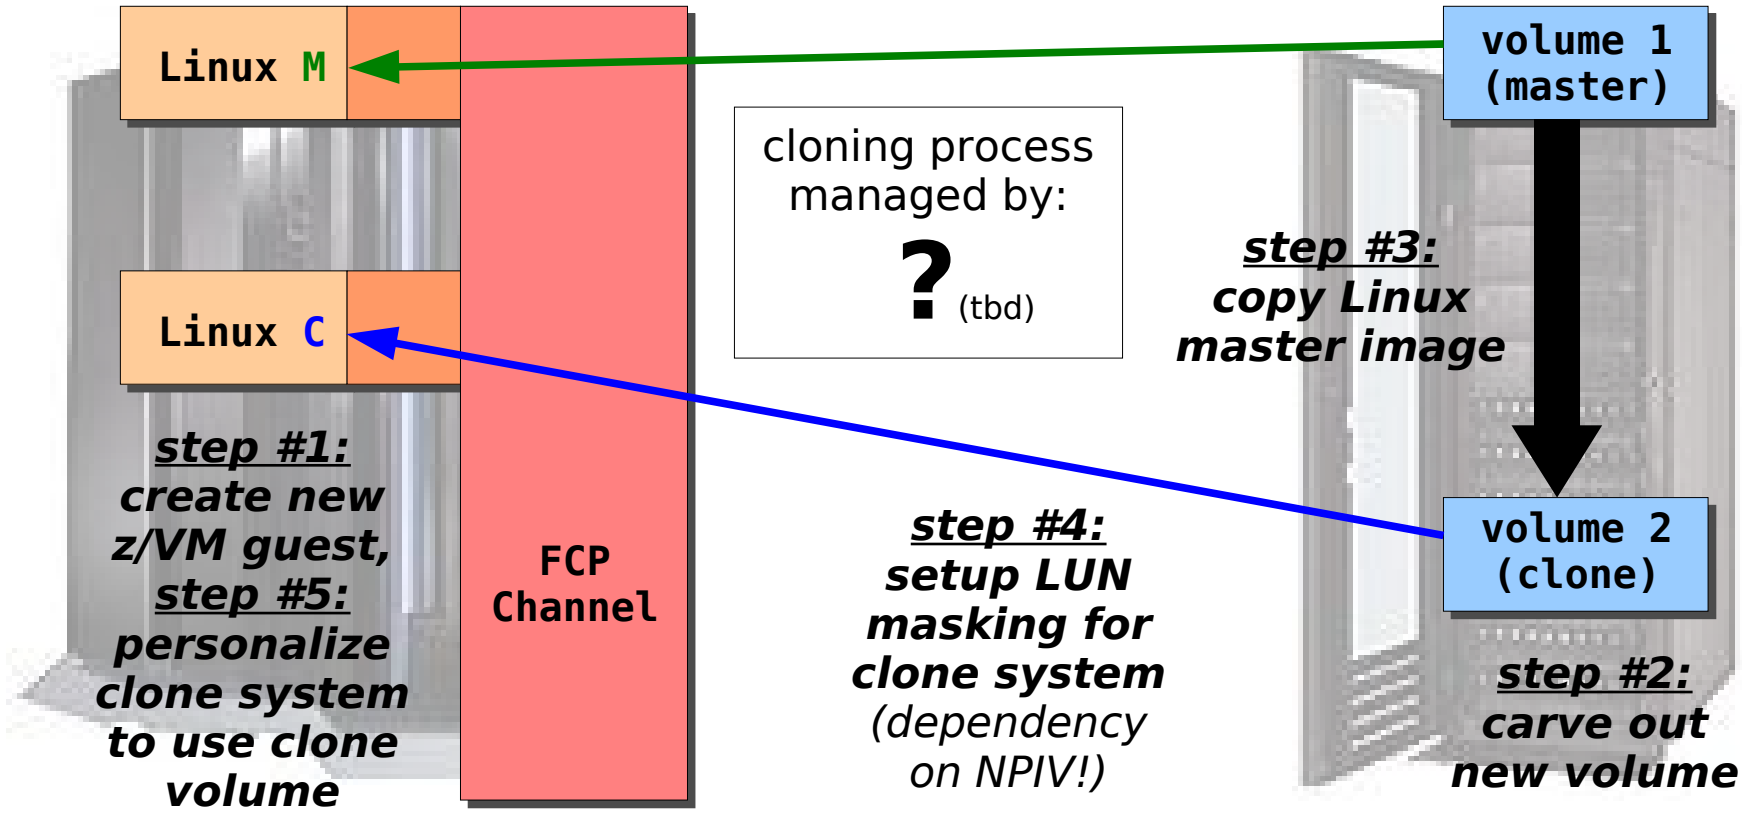

# Requirements


## IBM System z9

see also *IBM System z9 109*
http://www.ibm.com/servers/systems/systemz9/z9109/features.html

## NPIV-Capable Switch

- → only required for switch adjacent to z9
- → McData switches w/ firmware upgrade:

see also *E/OS 8.00*
http://www.mcdata.com/downloads/tc/cdesc/sphereon_4400_4700_wbt.pdf

# 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**
  (mostly conditions due to deficient planning)
- **Linux code to be shipped anytime soon**
  (target: SLES9 SP3 and equivalent)

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

## NPIV Step-by-Step

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

# 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 Fibre Channel Protocol for SCSI—FCP channels*
http://www.ibm.com/servers/eserver/zseries/connectivity/fcp.html

# Look'n Feel

## NPIV Is On

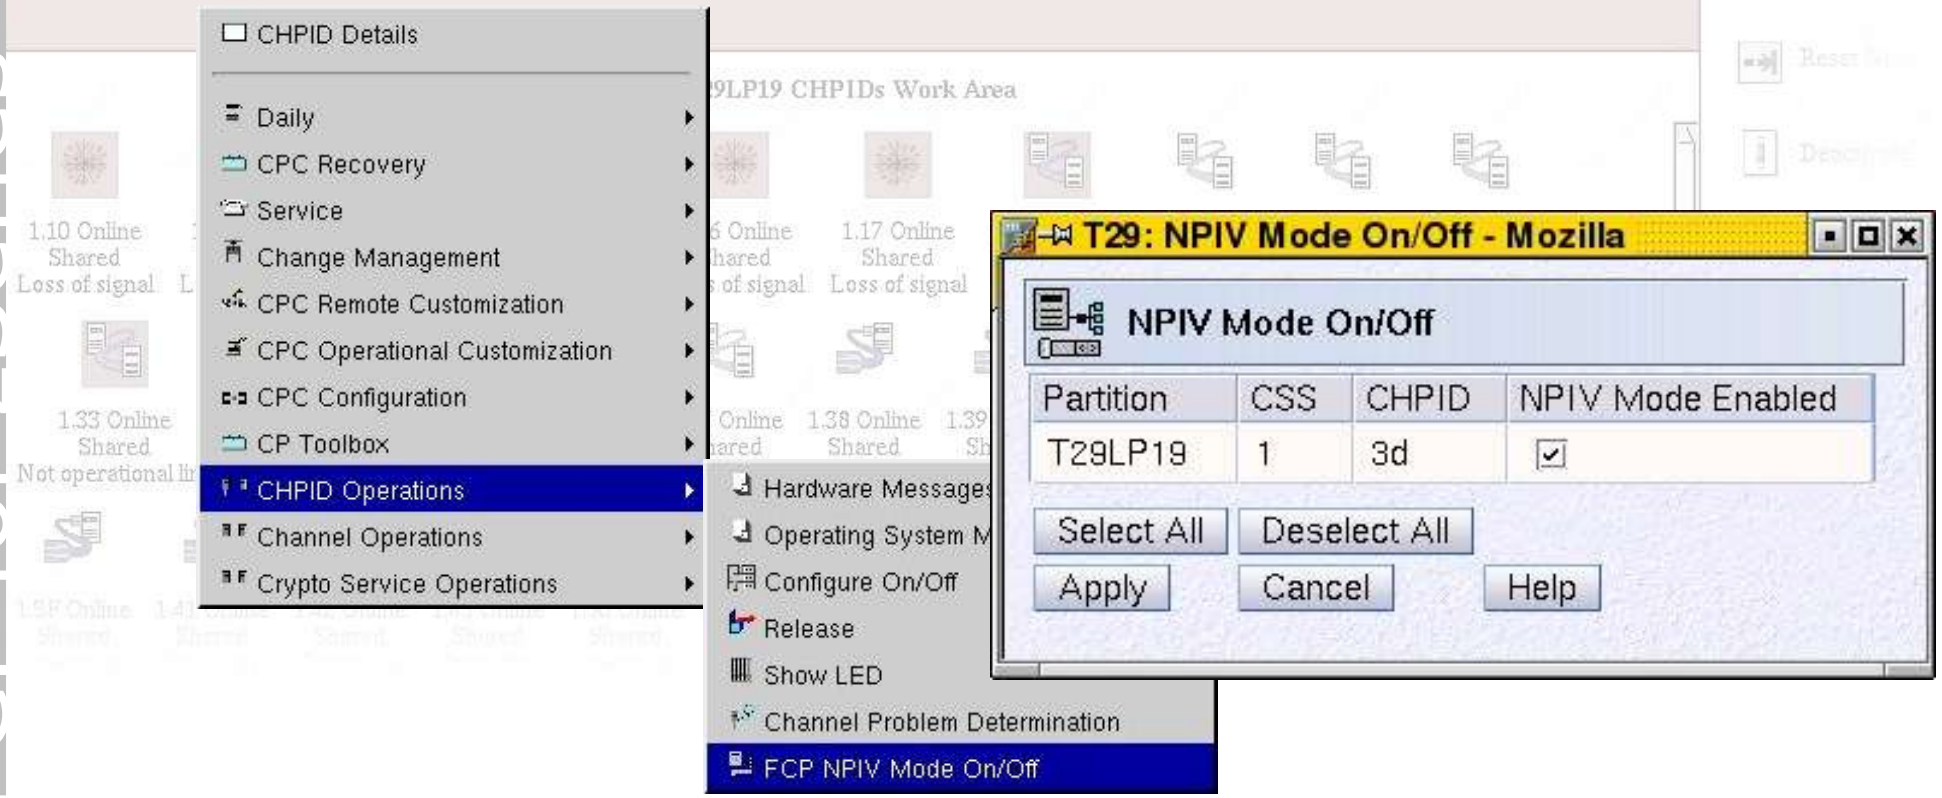


➔ **your choice: toggle NPIV on with a per channel-and-LPAR granularity**
(CHPID needs to be toggled off temporarily)

# Look'n Feel

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

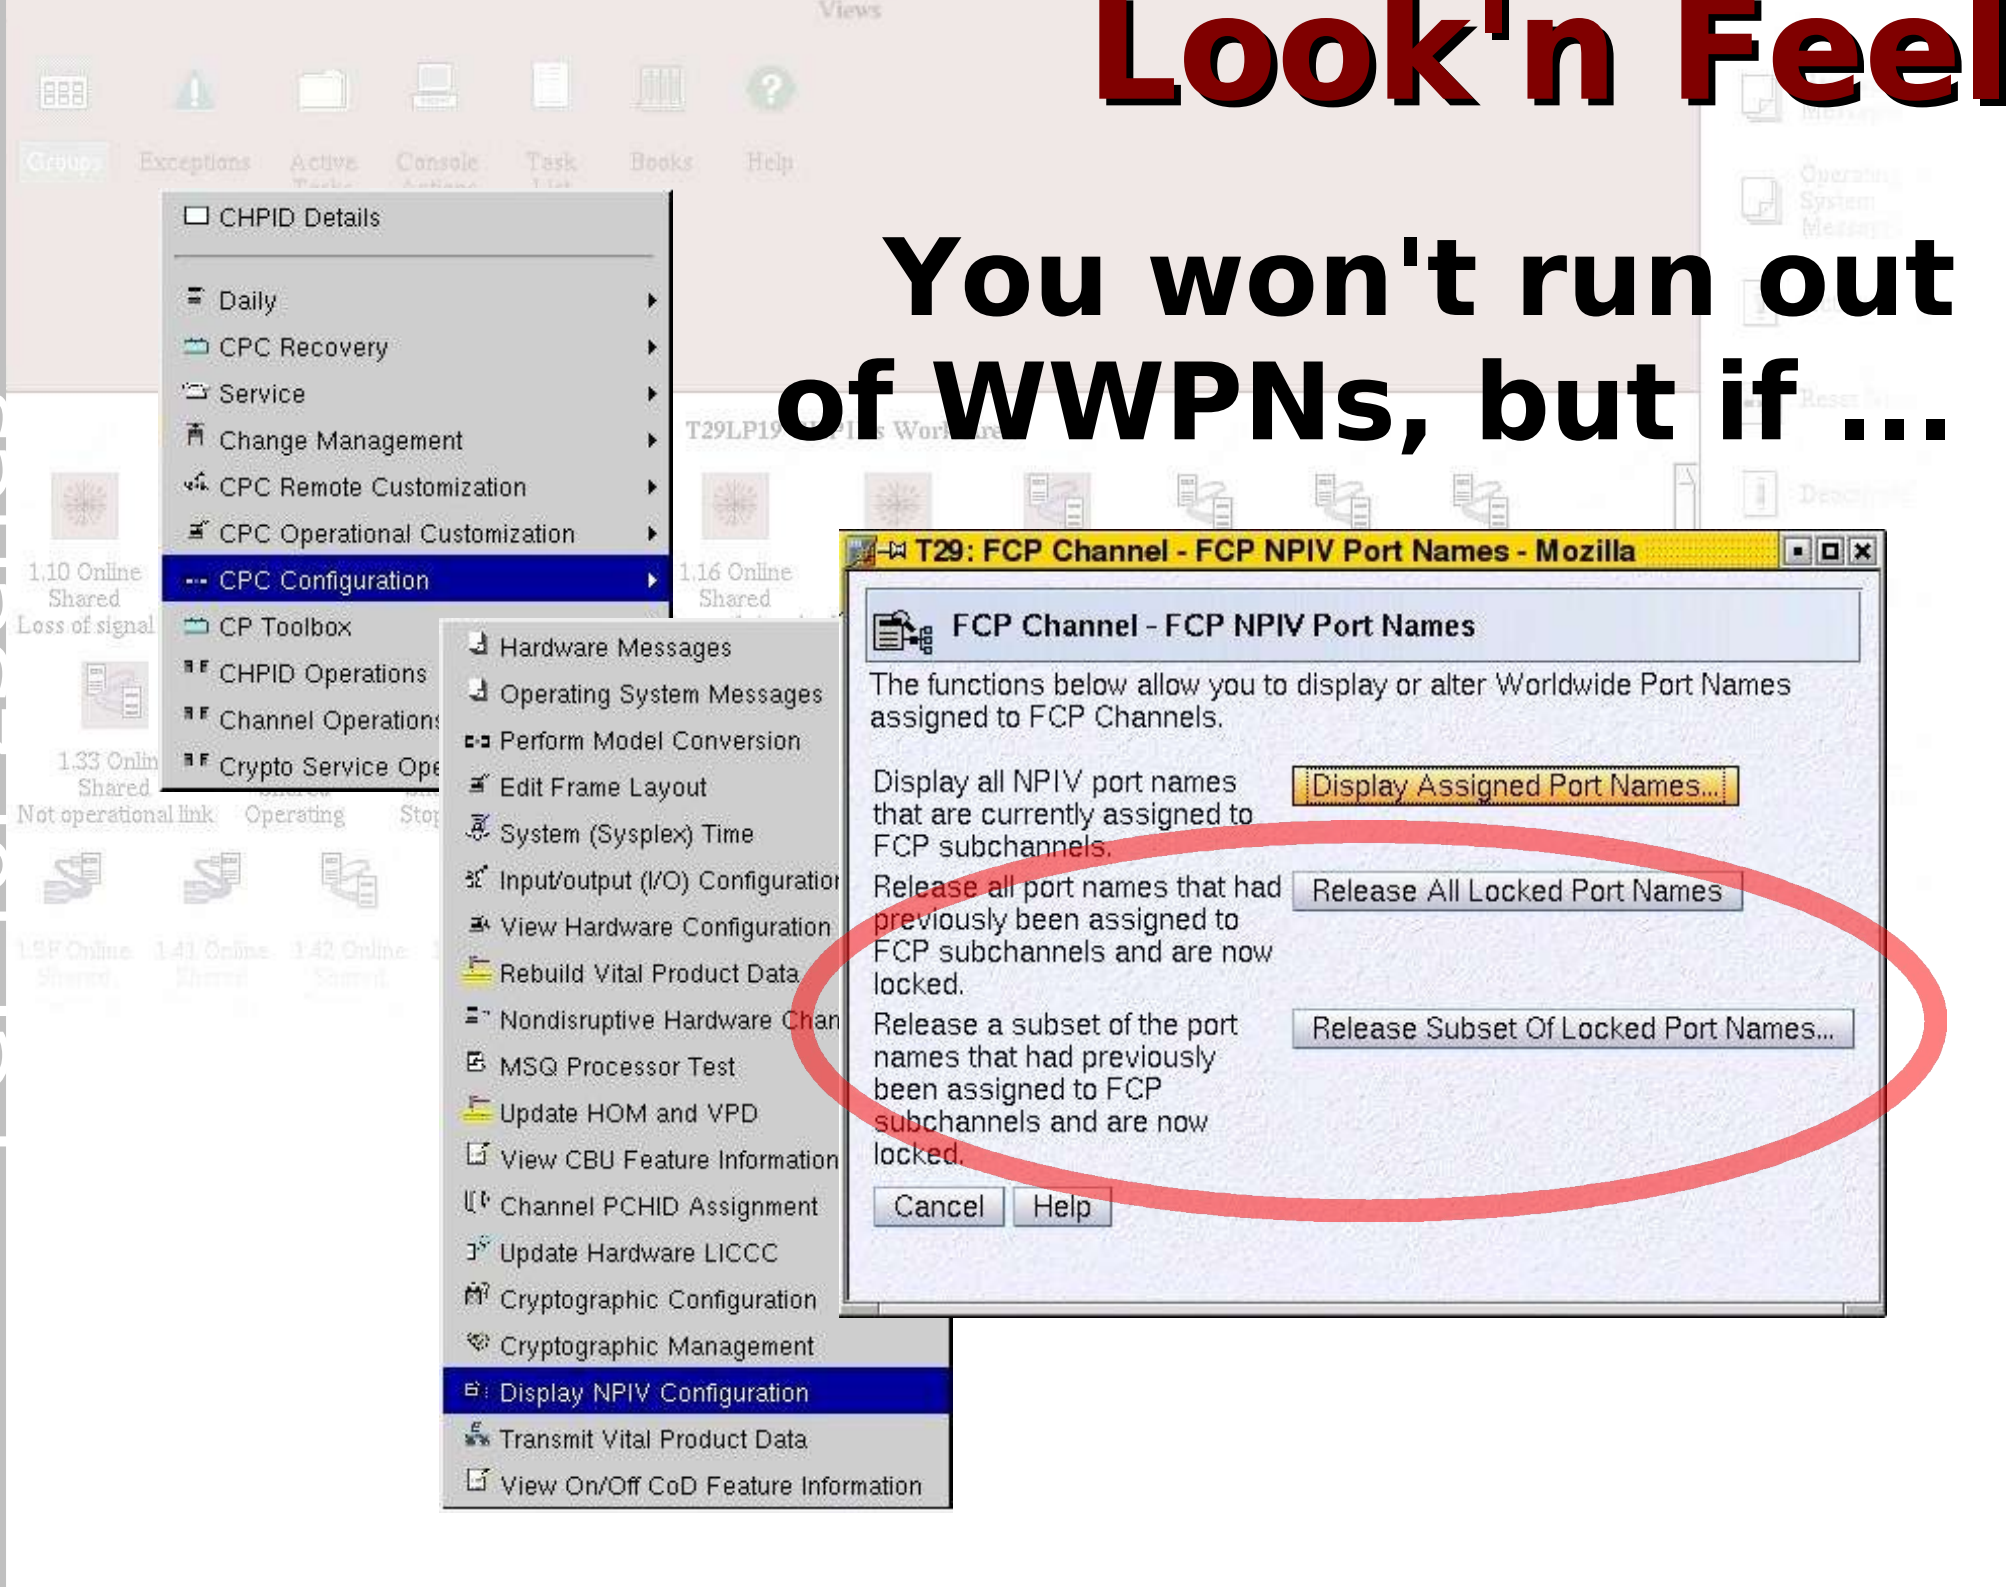

# Look'n Feel

## List Of Virtual N_Ports

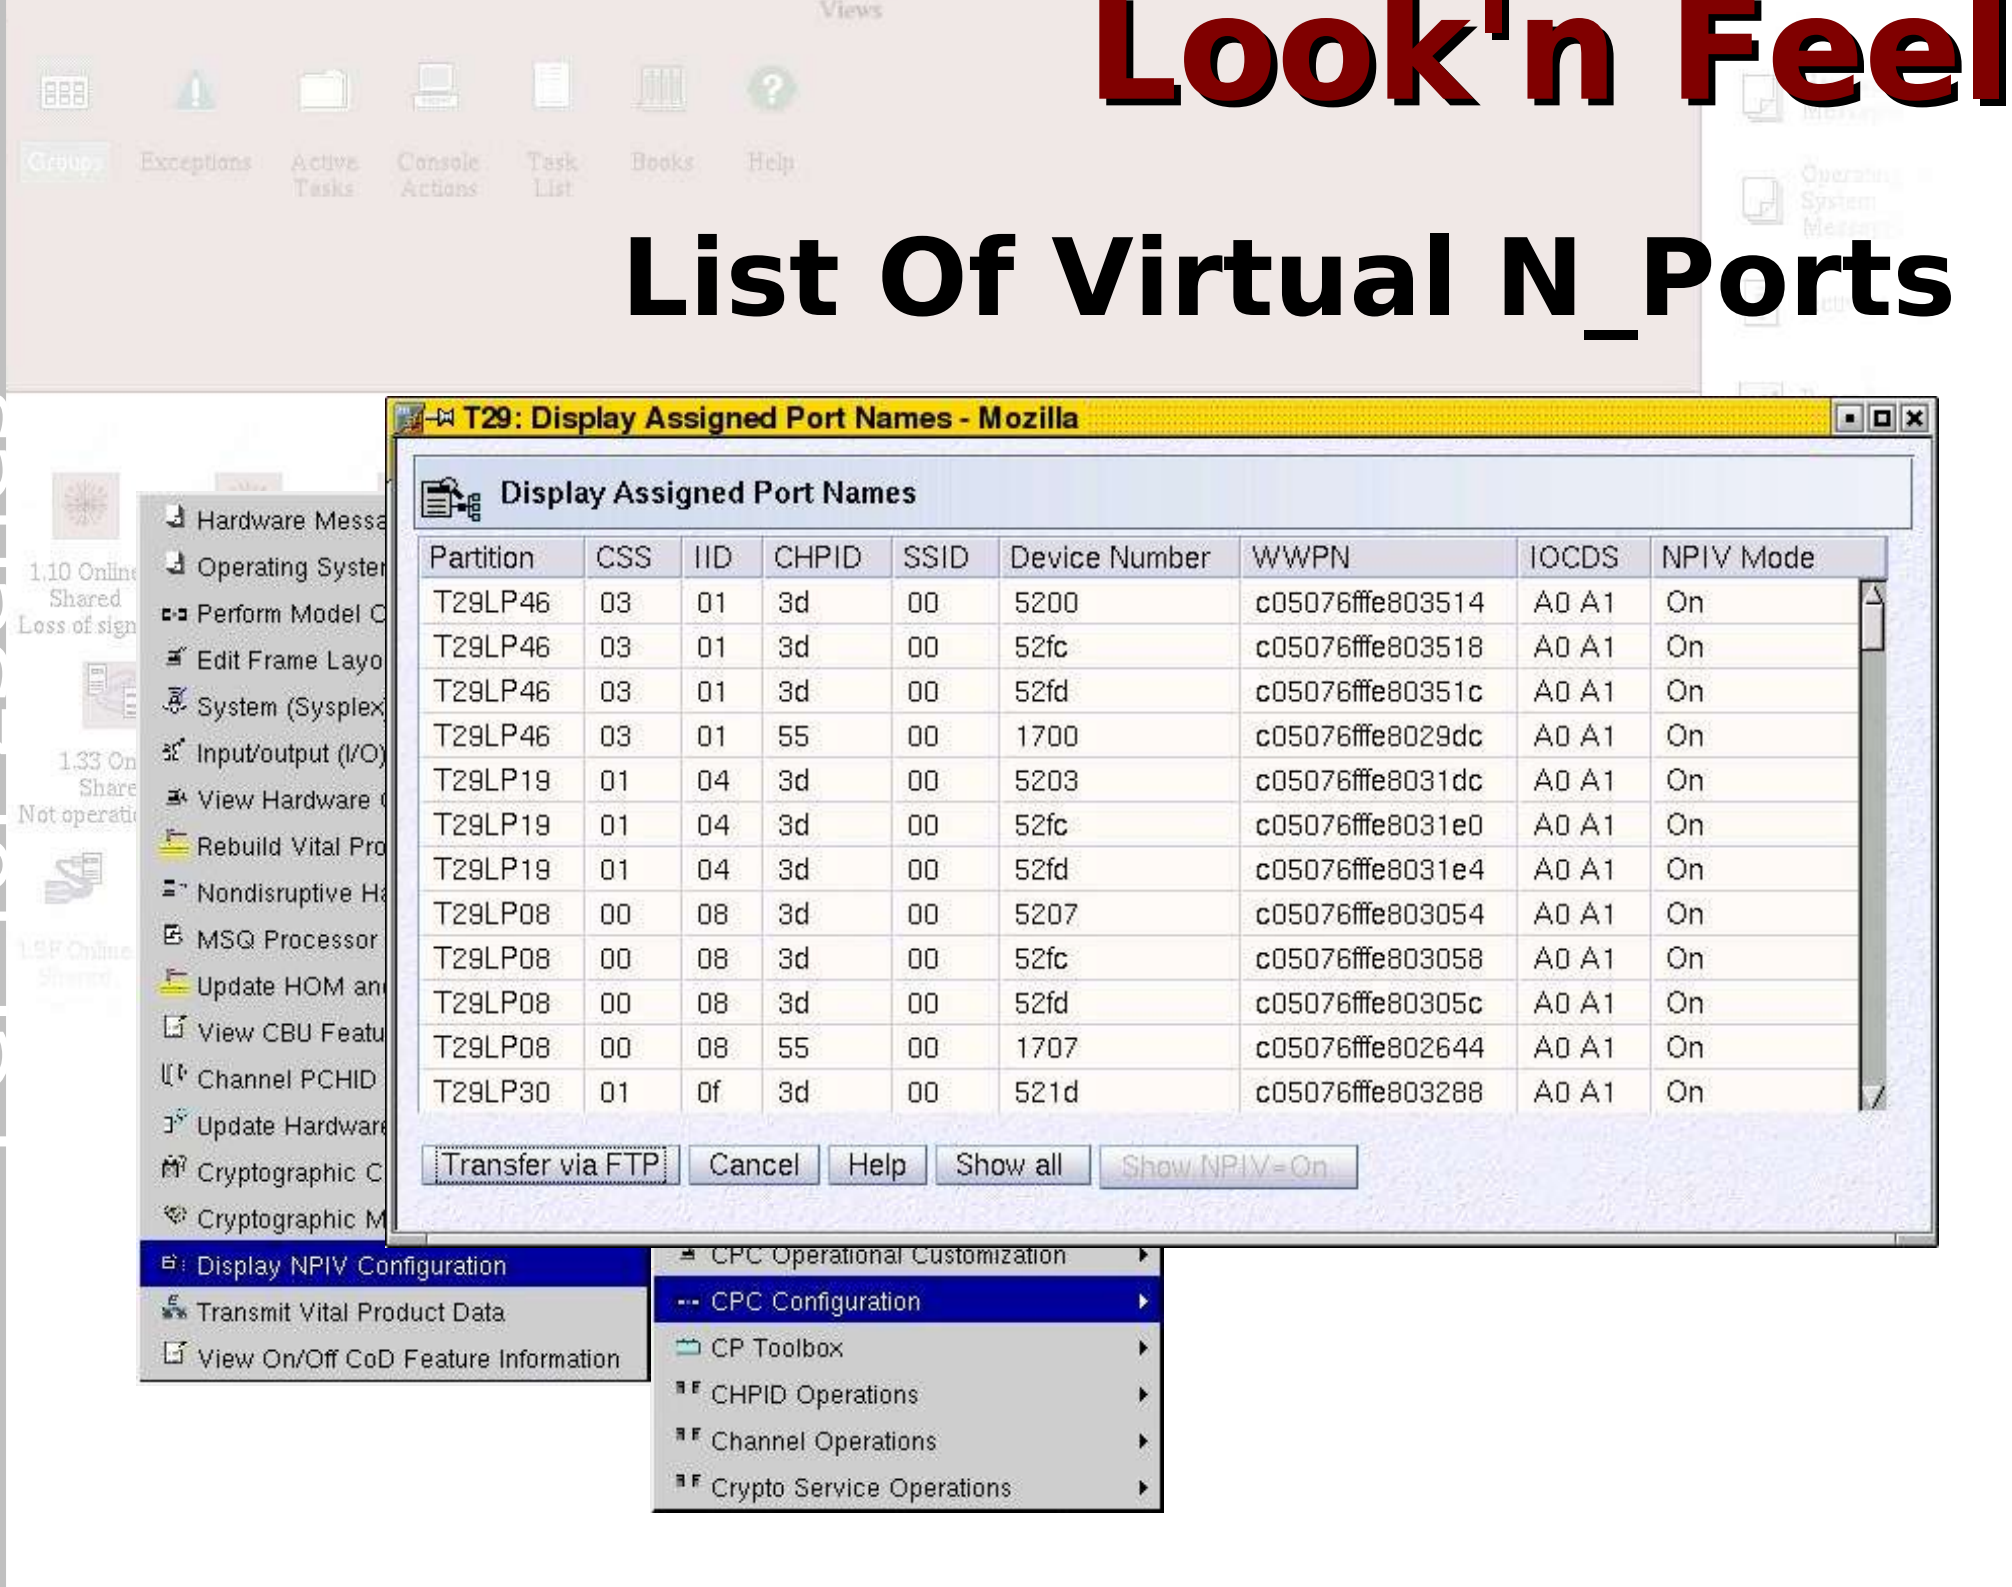

# Look'n Feel

## Linux Proudly Presents: Its Own SAN Identity

**FCP subchannel**

**virtual N_Port tied to this FCP subchannel**

*in NPIV mode:*
*values differ*
*in non-NPIV mode:*
*same WWPN, same S_ID*

**physical N_Port tied to FCP channel (CHPID)**

```
-bash-3.00# pwd
/sys/bus/ccw/drivers/zfcp/0.0.5203

-bash-3.00# cat wwpn s_id
0xc05076fffe8031dc
0x652214

-bash-3.00# cat physical_wwpn physical_s_id
0x5005076401e06b18
0x652213
```

# Summary

## Ready, steady, ... go!

- NPIV eradicates the root cause for limited practicability of FCP in virtual server environments
- NPIV enables the deployment of standard SAN management functions with System z9
- NPIV gives free rein to sophisticated SAN fantasies