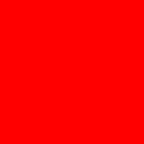


**ORACLE®**



The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.



**ORACLE®**

## **State of the LDom Version 2.1**

Stefan Hinker  
EMEA Hardware Principal Sales Consultant

**ORACLE®**

**VM**

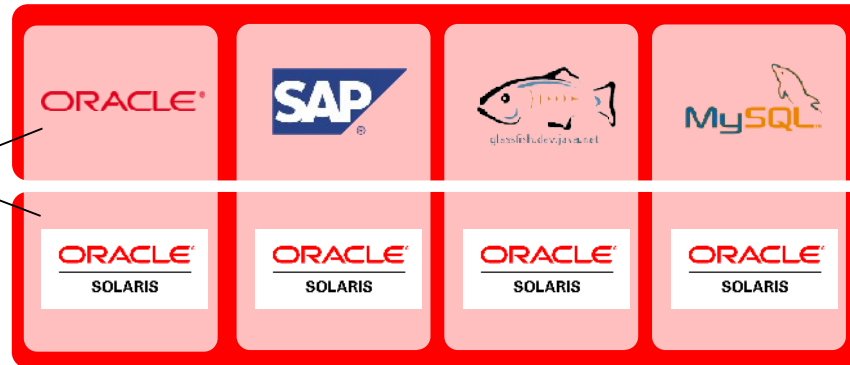
# Agenda

- **Very Short Introduction to LDOMs**
- What's New in Version 2.1
  - Secure Live Migration
  - Dynamic Reconfiguration & Power Management
  - Inter-Vnet LDC Channels
- Management with OpsCenter
- LDOMs in SuperCluster
- On or Behind the Horizon...




# Oracle VM Server for SPARC

Isolated OS and applications in each logical (or virtual) domain



Firmware-based hypervisor



Each logical domain runs in dedicated CPU thread(s)

## Optimized for SPARC / Oracle Solaris

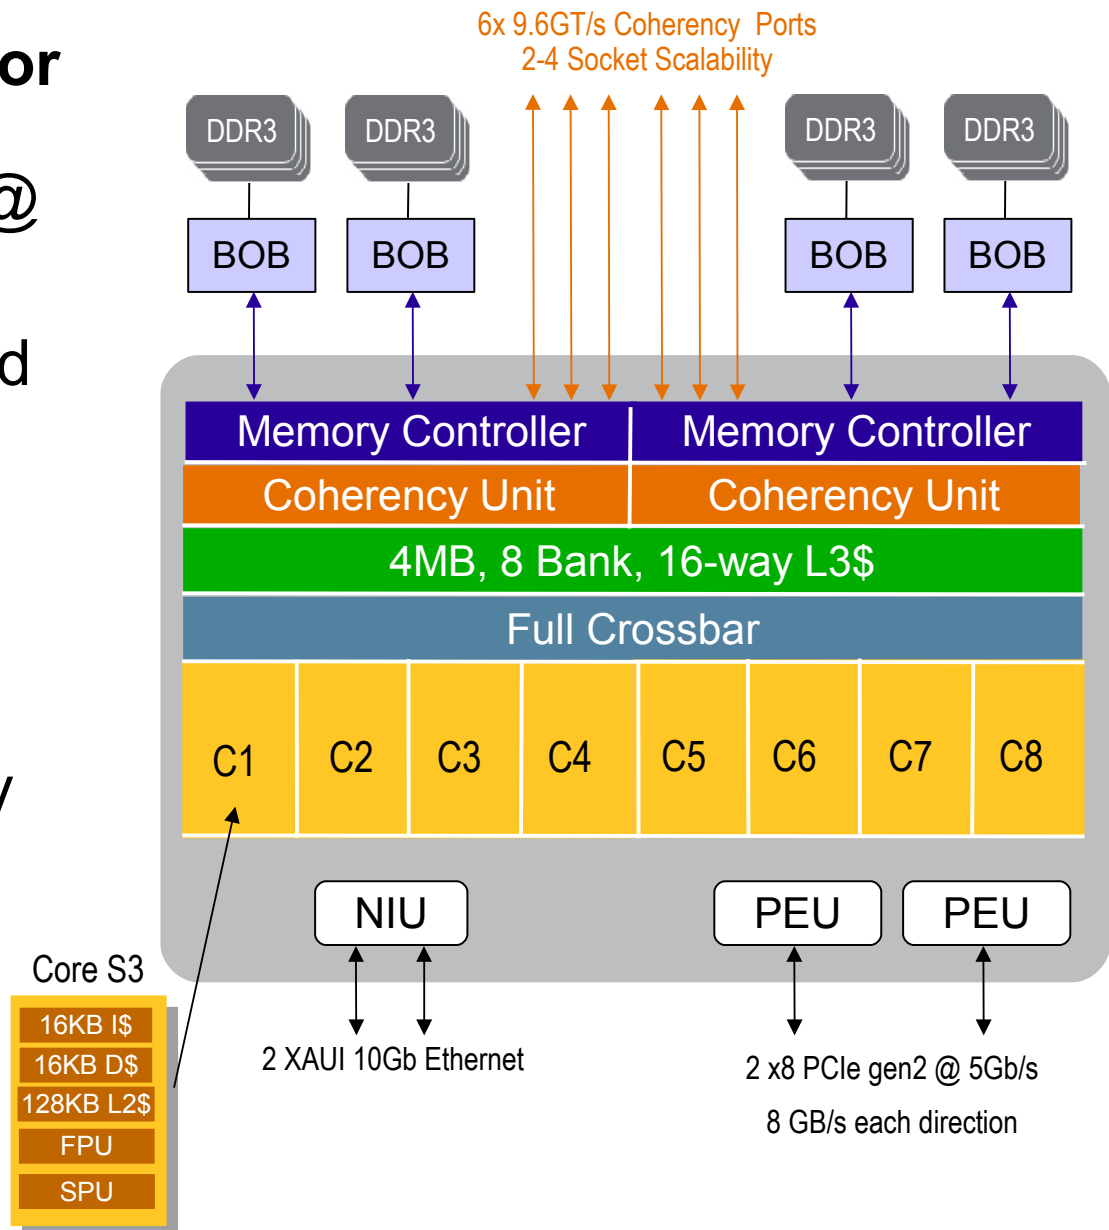
# OVM Server for SPARC 2.1 Prerequisites

- UltraSPARC T2(+), SPARC T3 or T4
  - No support for T1!
- Firmware
  - 7.4 for T2/T2+
  - 8.1 for T3 & T4
- Operating system releases
  - Solaris 10 9/10 (Update 9)
  - Solaris 11
  - Release Notes contain a list of required patches for older releases

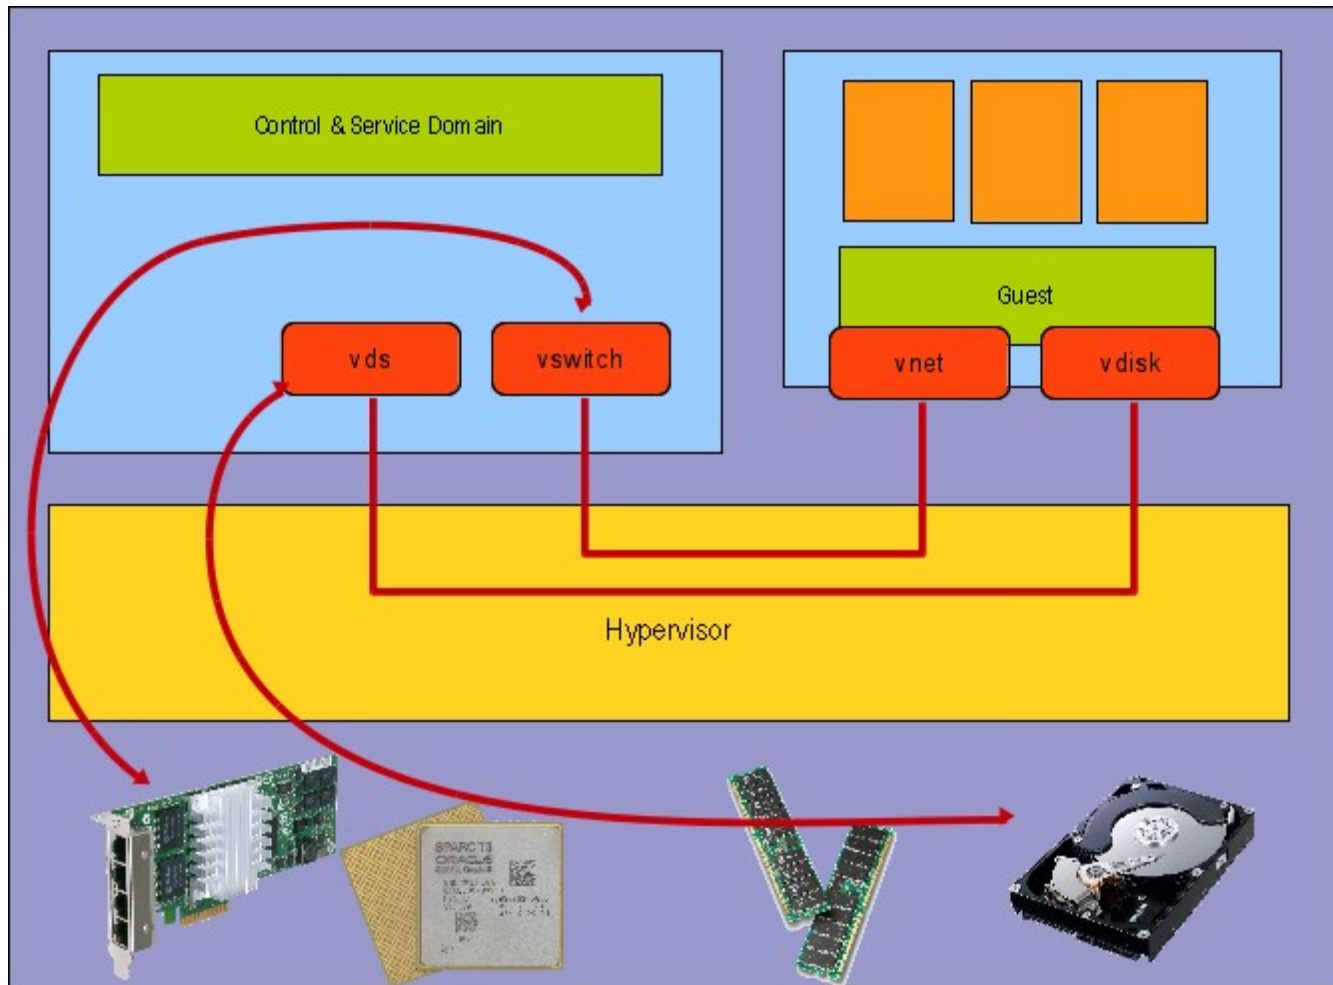


# SPARC T4 Processor

- 8 cores, 8-64 threads @ up to 3Ghz
- Single or multi-threaded operation per core
- System scalability to 4 sockets
- SPARC Core “S3”
- 1-8 Strand Dynamically Threaded Pipeline
- ISA-based Crypto-acceleration
- 4MB Shared L3\$

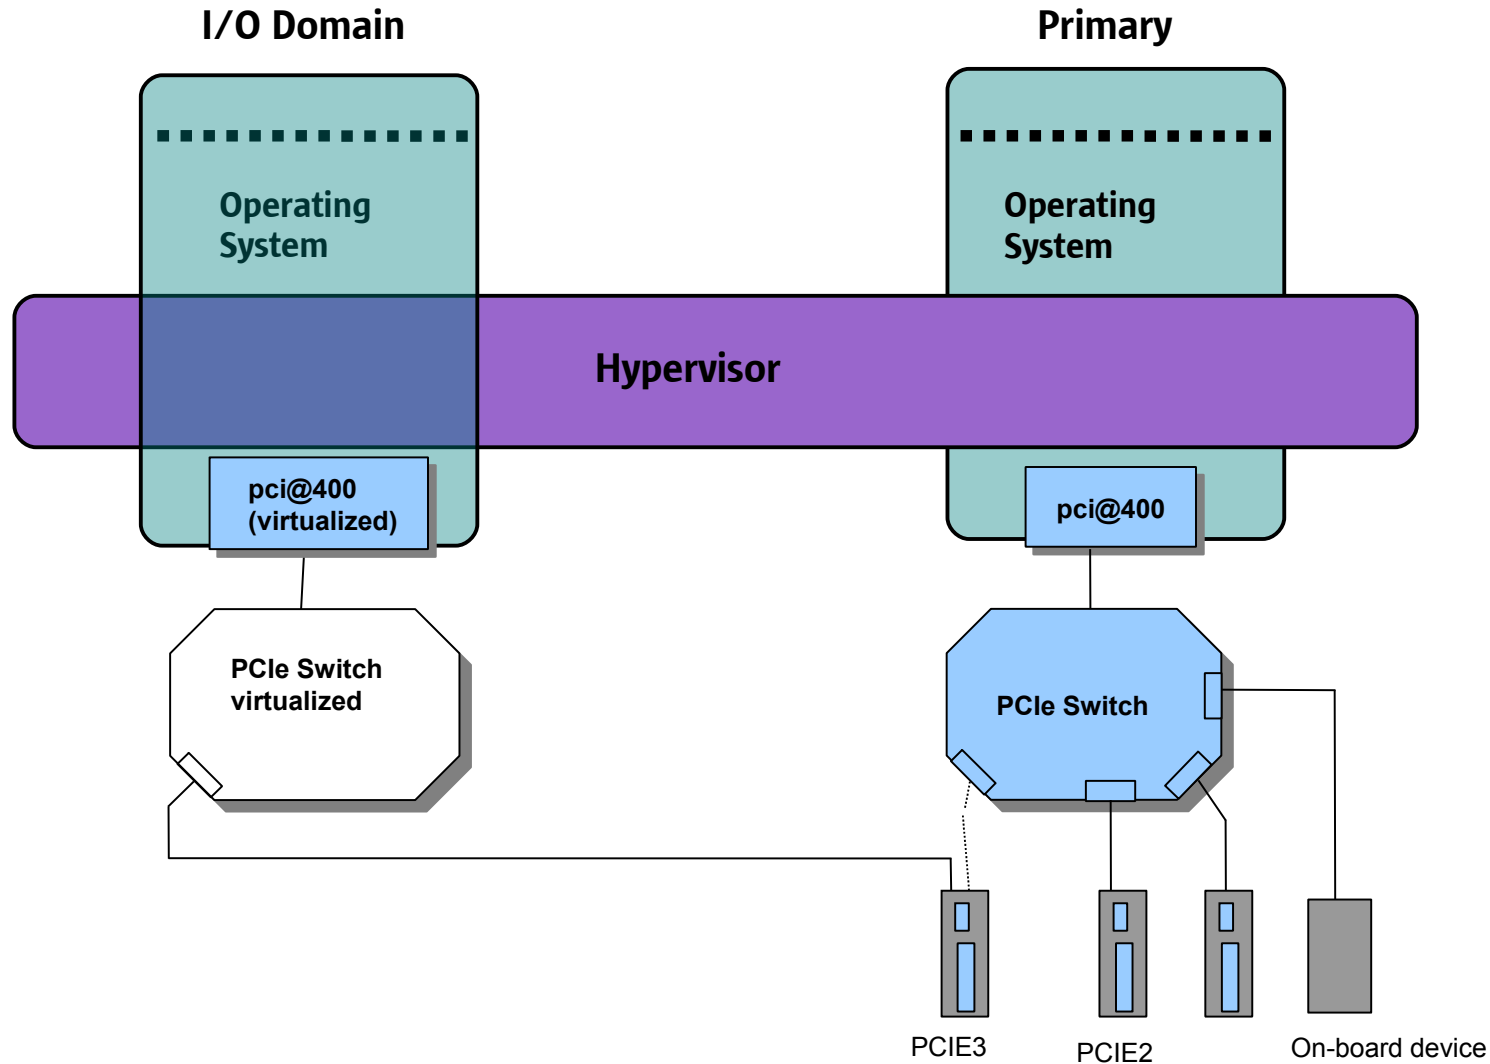


# Domain Components

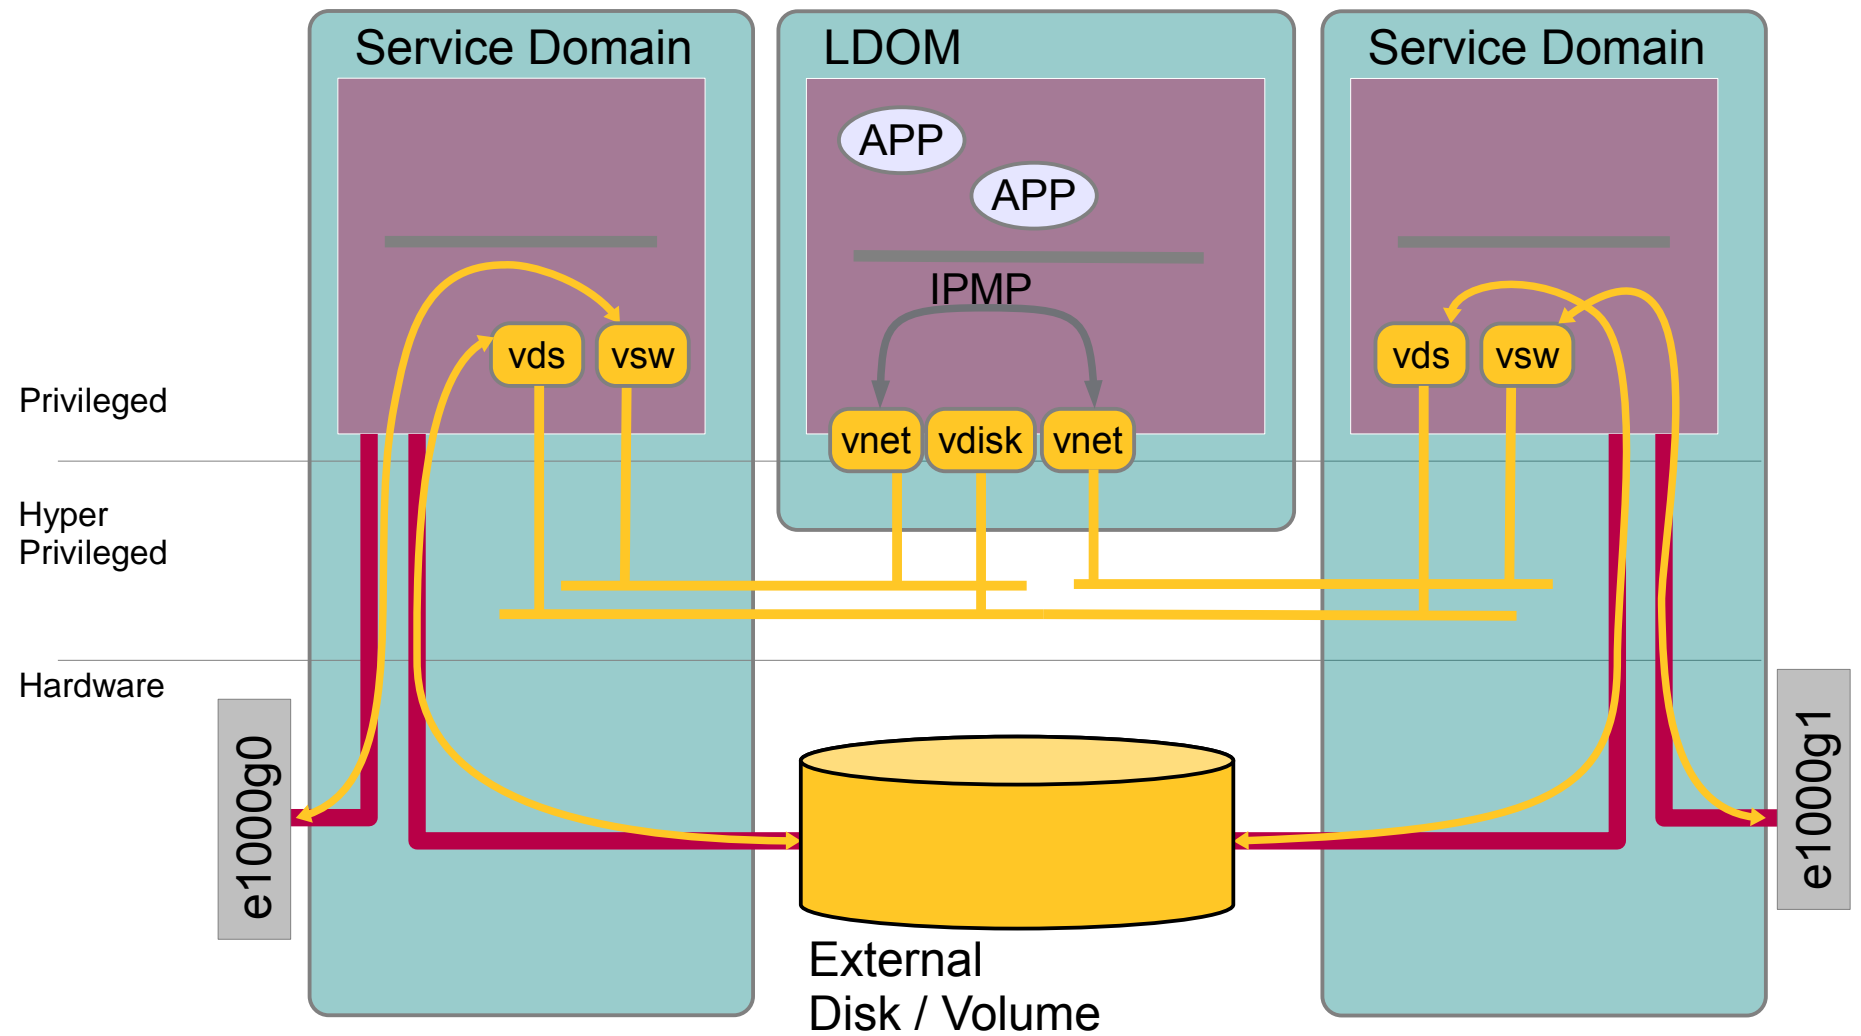




# Static Direct IO (SDIO)



# Redundant Virtual I/O



# Agenda

- Very Short Introduction to LDoms
- **What's New in Version 2.1**
  - Secure Live Migration
  - Dynamic Reconfiguration & Power Management
  - Inter-Vnet LDC Channels
- Management with OpsCenter
- LDoms in SuperCluster
- On or Behind the Horizon...



# Oracle VM Server for SPARC 2.1

## How did we get here ?

- LDoms 1.0: 2007 on T1000 and T2000
- LDoms 1.1: 2008 on T2/T2+
- LDoms 1.2: 2009
- OVM Server for SPARC 2.0 1.3: 2010
- OVM Server for SPARC 2.0: 2011 on T3
- OVM Server for SPARC 2.1: 2012 on T4

<http://www.oracle.com/technetwork/server-storage/vm/documentation/sparc-whatsnew-330281.html>

# LDoms – What's New in Release 1.1

- Guest Domain Migration
- Dynamic Reconfiguration for Virtual IO
- NIU Hybrid IO – Direct Access to 10GBit Ethernet
- Virtual Disk Multipathing
- Support for VLAN tagging

# LDoms – What's New in Release 1.2

- Support for Jumbo Frames
- Support for Multi-Initiator virtual Disks
- Less “delayed reconfiguration”
- Customizable Domain dependencies
  - Support for Solaris Cluster
- P2V Tool
- Configuration Assistant
- CPU Power Management

# OVM Server for SPARC – What's New in Release 1.3

- Link-based IPMP
- Migration of Domains with Crypto-Units
- Dynamic CPU Resource Management
- hostid and MAC freely configurable
- Faster migration through memory compression
- Bootdevices > 1TB

# OVM Server for SPARC – What's New in Release 2.0

- Memory DR
- Static Direct IO
- Cooperative Guest Domain Migration
- Whole Core allocation & affinity binding
- Memory Power Management & Power Capping
- Virtual Disk Multipathing
- LDoms Agent



# OVM Server for SPARC – What's New in Release 2.1

- Secure Live Migration
- Dynamic Resource Management Enhancements
- CPU Affinity
- P2V Enhancements
- Inter–Vnet LDC Channels
- Virtual Device Service Validation
- Updated LDoms MIB
- Automatic Crypto Removal During DR
- Other RFEs implemented

# Agenda

- Very Short Introduction to LDoms
- What's New in Version 2.1
  - **Secure Live Migration**
  - Dynamic Reconfiguration & Power Management
  - Inter-Vnet LDC Channels
- Management with OpsCenter
- LDoms in SuperCluster
- On or Behind the Horizon...



# Domain Migration Definitions

- Migration
  - move a domain from one system to another
  - This is not cloning, HA or replication for disaster recovery
    - use 'ldm ls-constraints -x <ldom>'
- Cold Migration: domain is inactive or bound
- Live Migration: domain is active
  - very limited suspend time
  - state transfer begins before suspending the domain

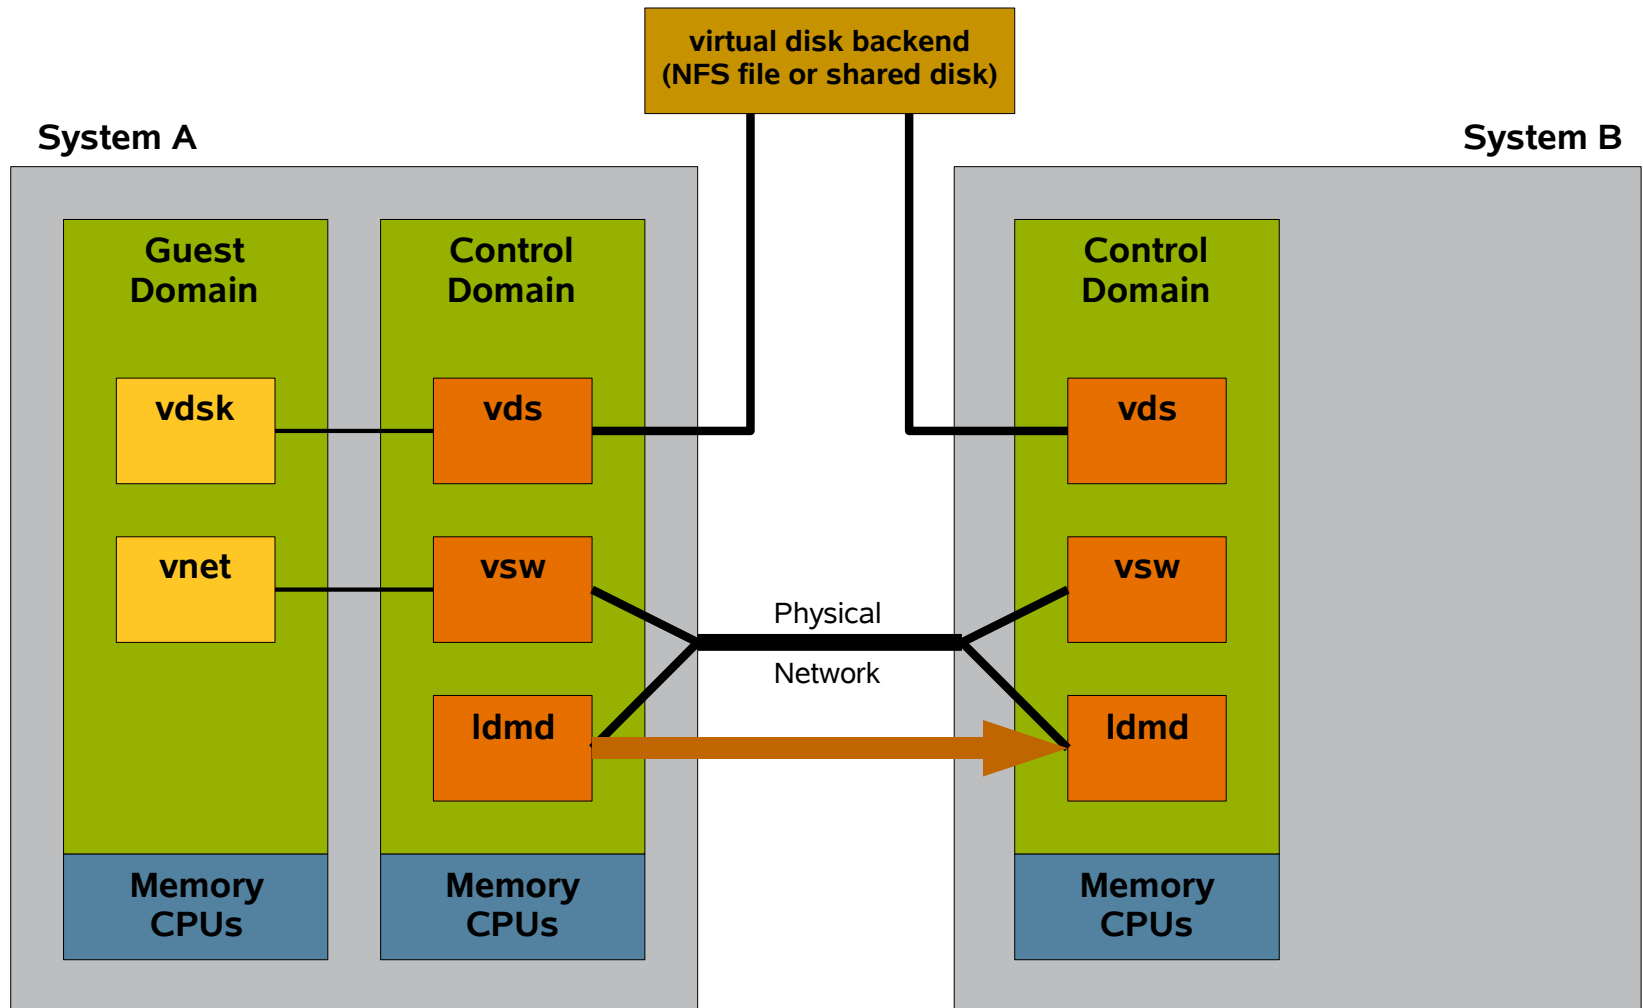
# Domain Migration with 2.1

- Live Migration - #1 feature request from customers
- Some Requirements for use
  - Source and Target systems must be running LDoms 2.1 and updated firmware.
  - SysFW 7.4 or 8.1 requires 2.1 for live migration (but 2.1 will run on down-rev SysFW)
  - Same CPU Type & Frequency on source and target (for now)

# Domain Migration CLI

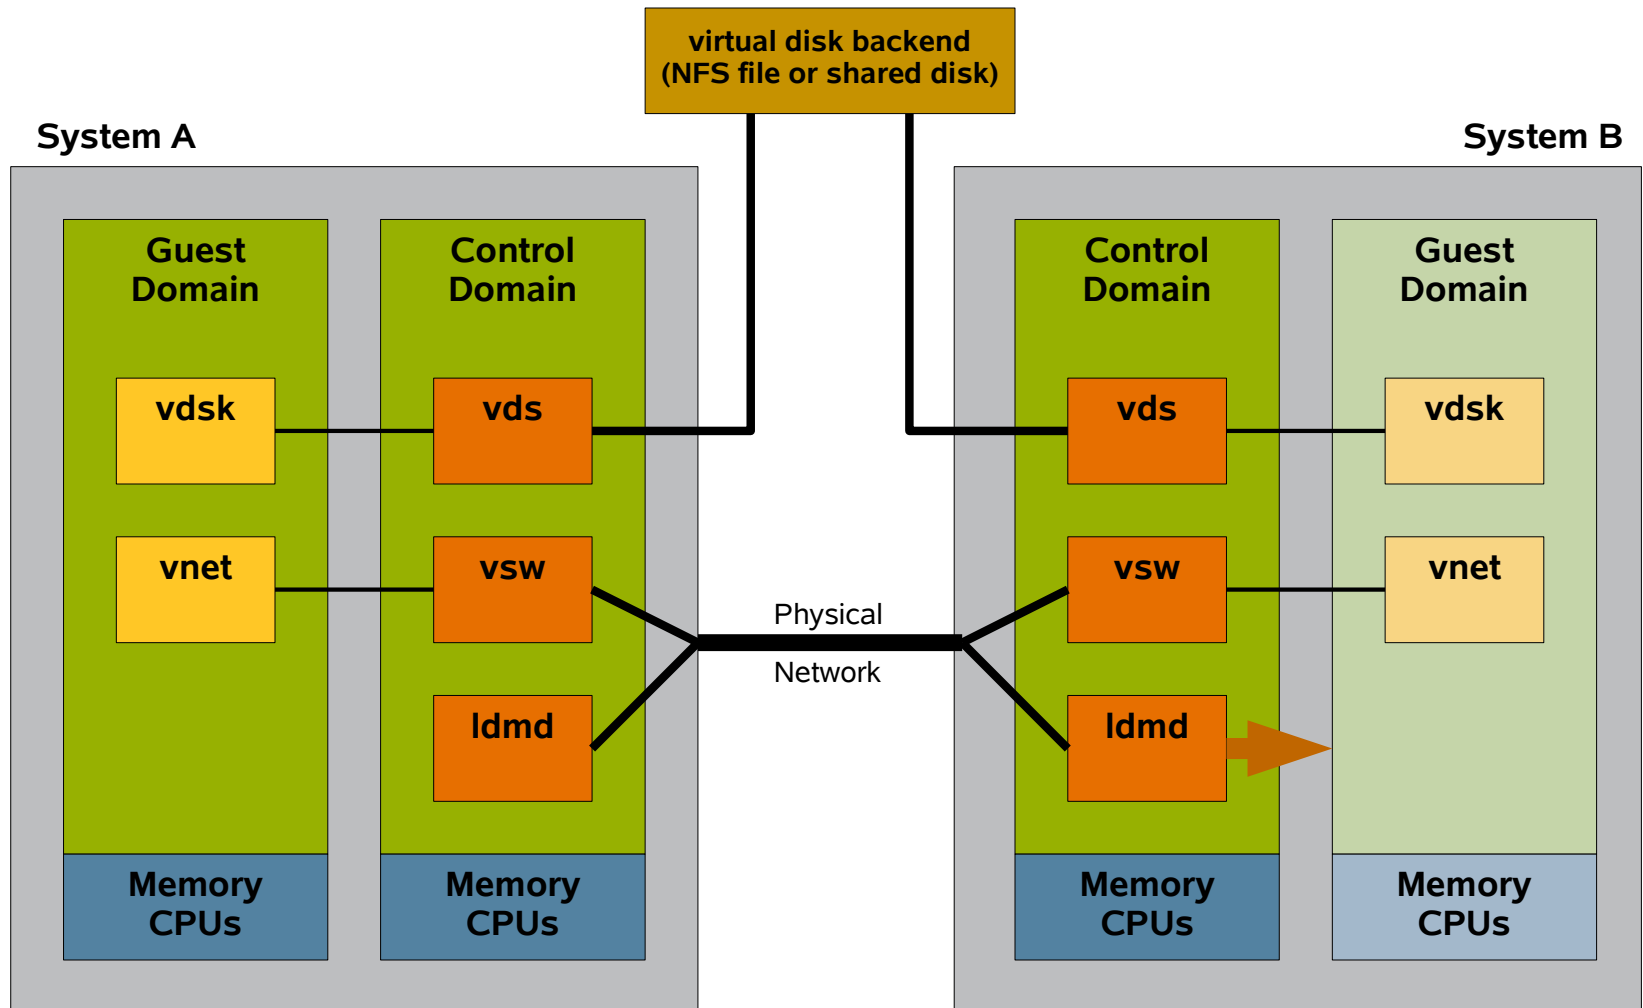
- Uses same CLI and XML interfaces as in prior releases
- To Migrate a Domain
  - `Idm migrate [-f] [-n] [-p <password_file>] <source-ldom> [<user>@]<target-host>[:<target-ldom>]`
    - -n : dry-run option
    - -f : force
    - -p : specify password file for non-interactive migration
- Cancel an On-Going Migration
  - `Idm cancel-operation migration <ldom>`
- Same Command for Cold and Live Migration
  - type of migration depends on the state of the domain

# Live Migration 1 of 6 - Initialisation



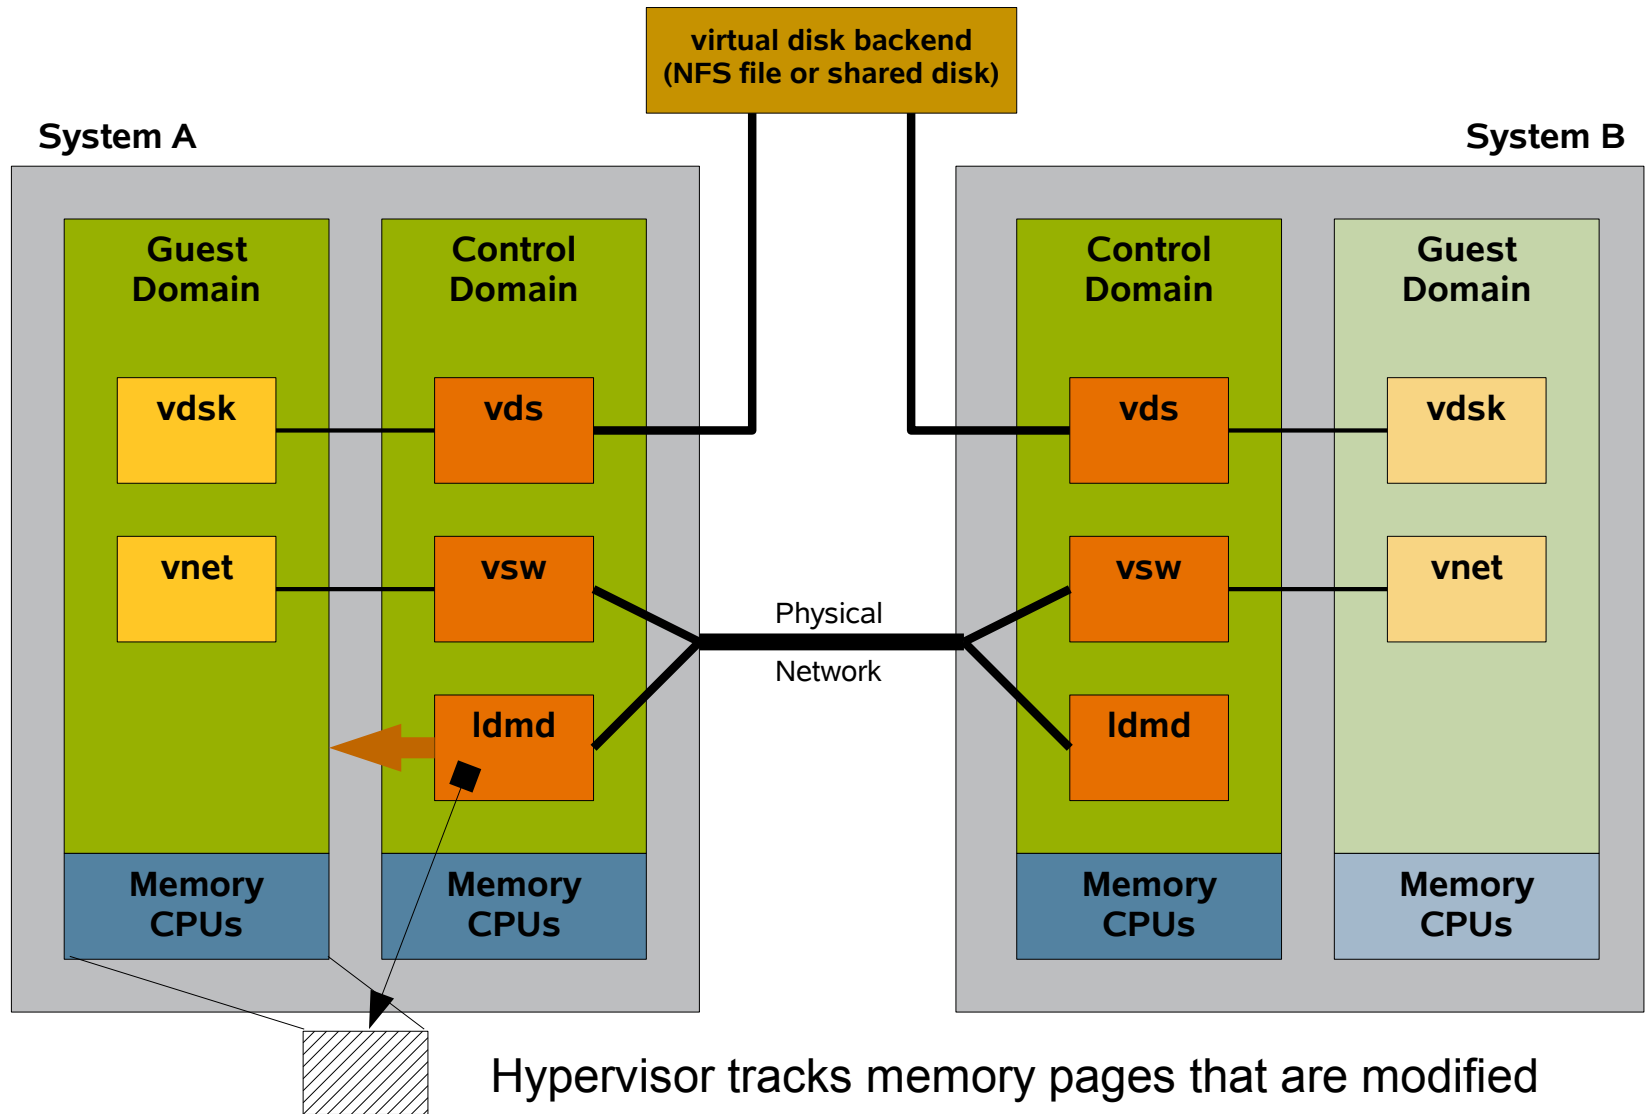
ldmd A checks with ldmd B if migration is possible, issues warning messages if not.

# Live Migration 2 of 6 – Create Target



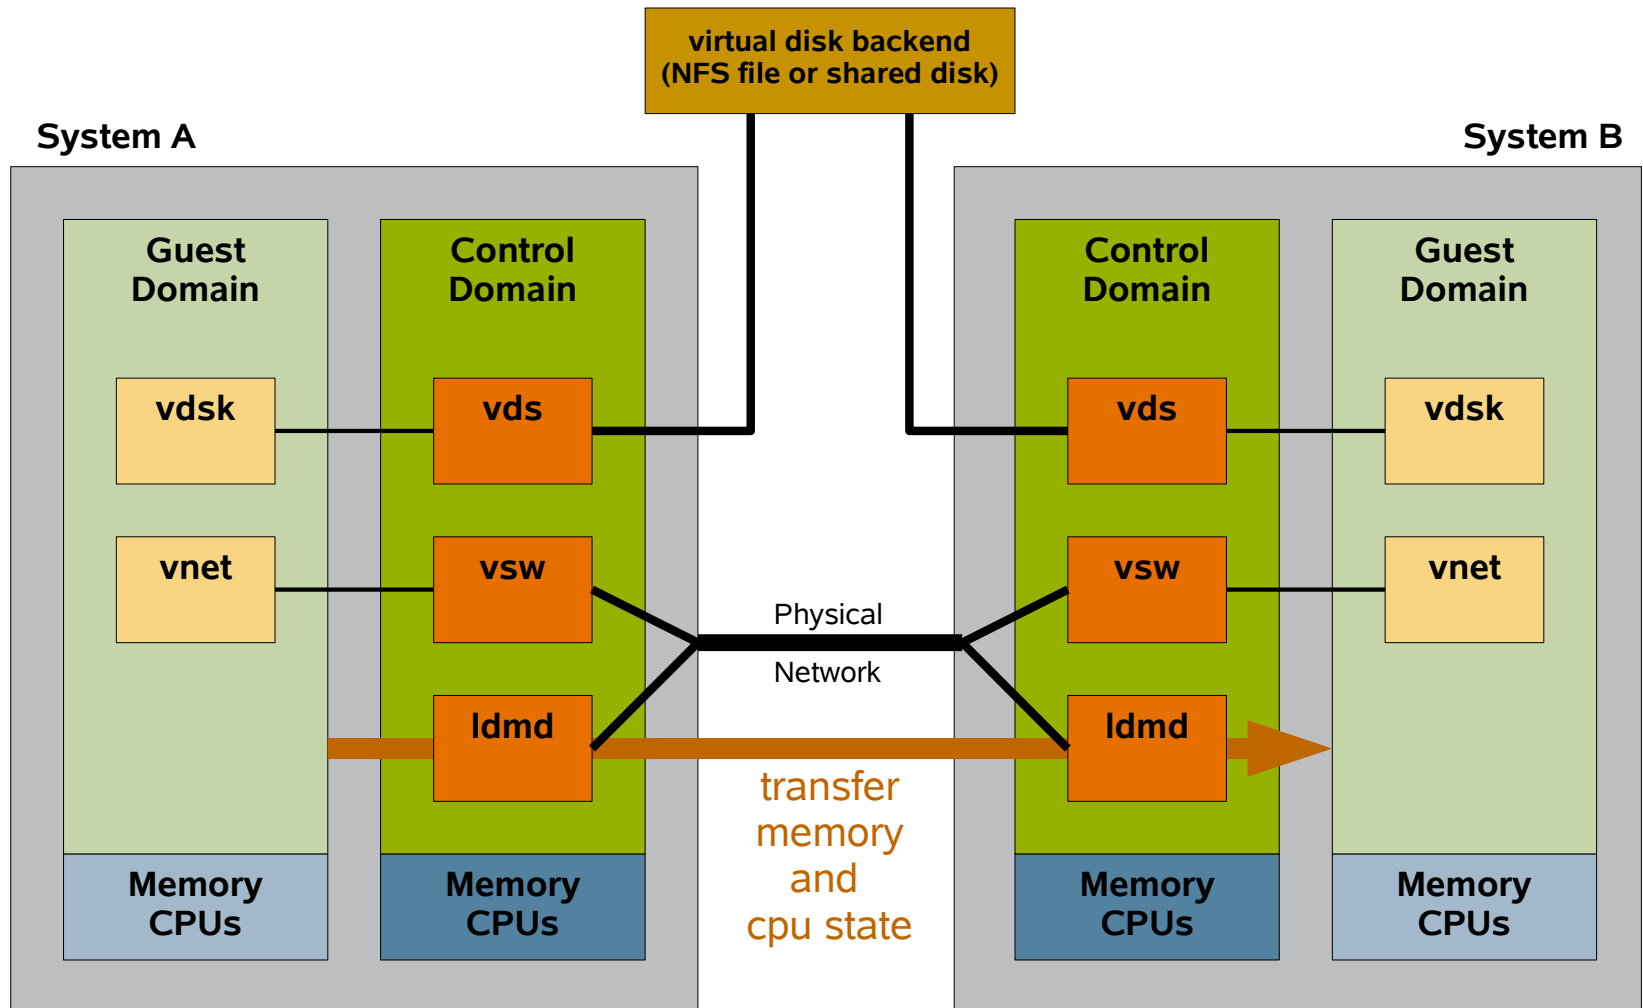
Idmd B creates and binds a domain with same amount of CPUs/Memory

# Live Migration 3 of 6 – Track Memory



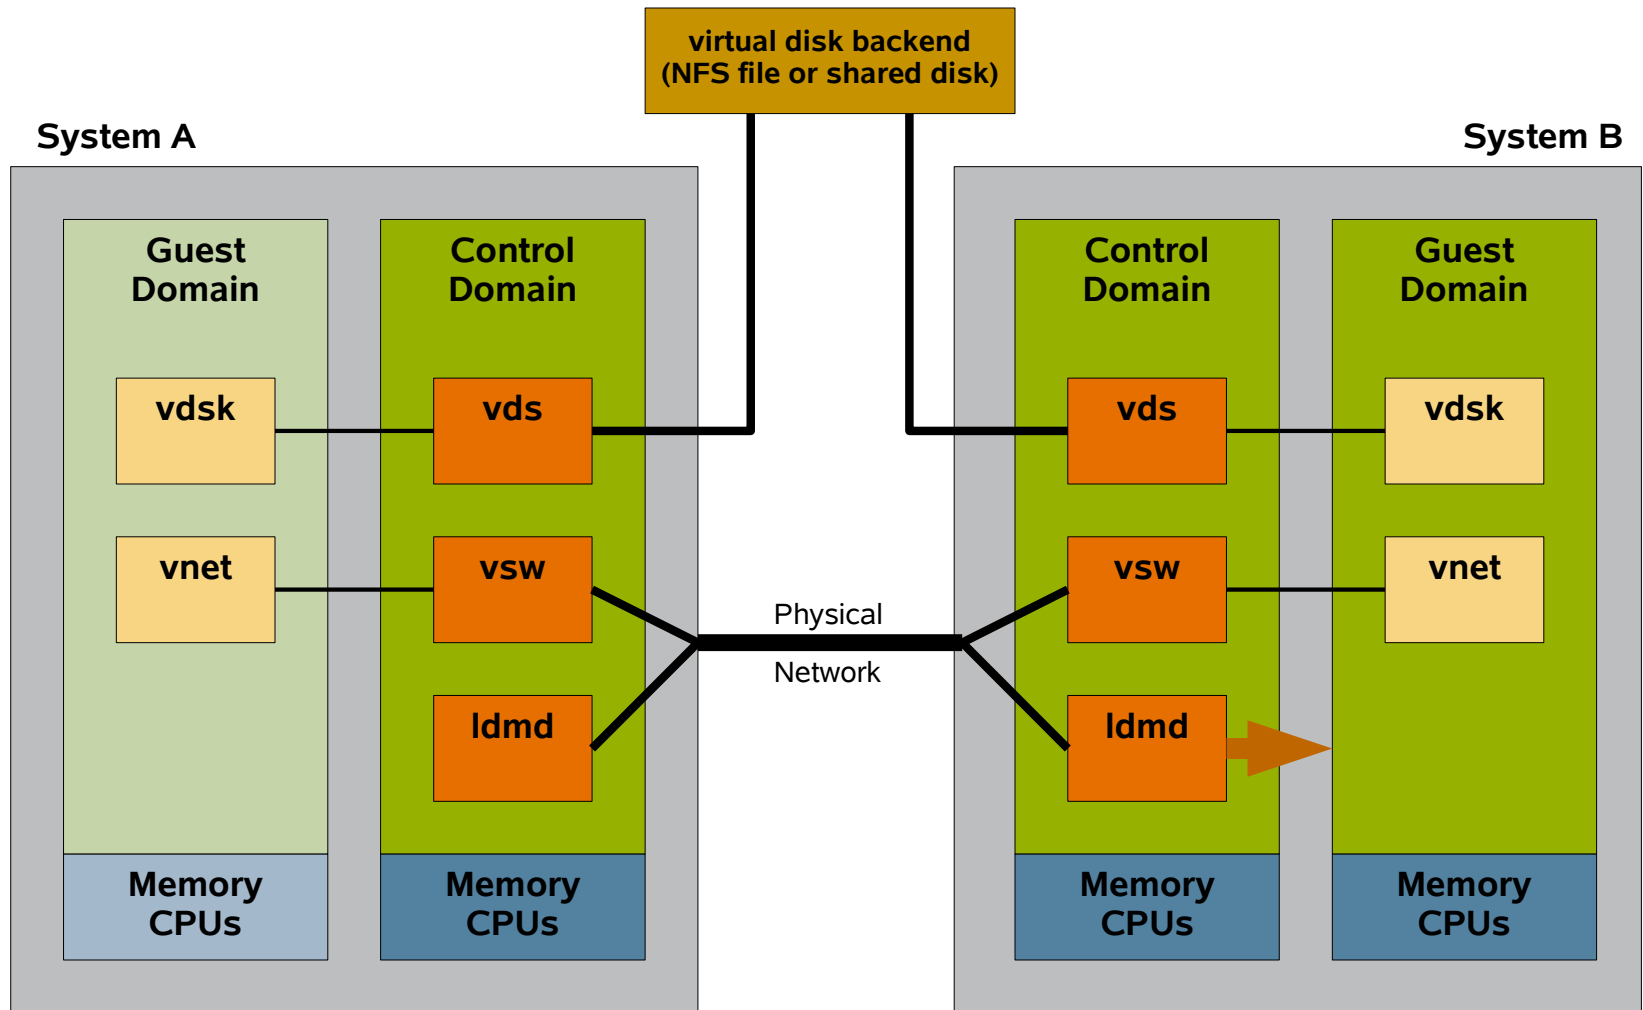


# Live Migration 4 of 6 – Transfer State



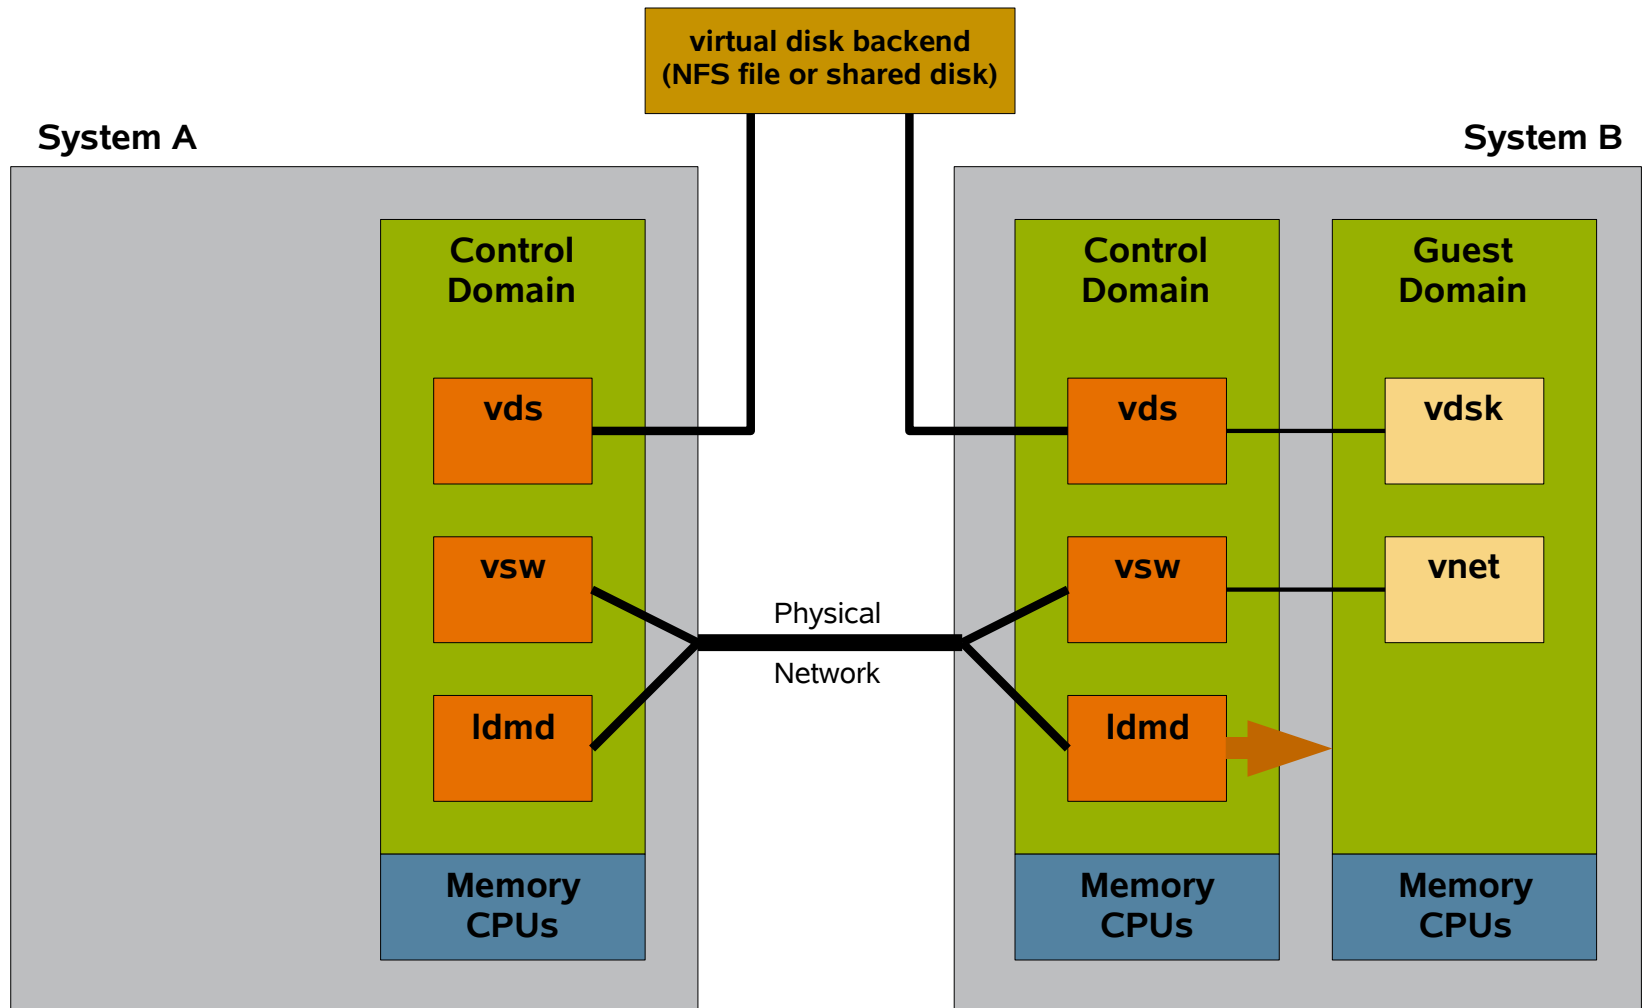
Idmd A transfers state while domain active and after multiple passes, suspends domain

# Live Migration 5 of 6 – Resume Target



Idmd B resumes the domain on the target machine

# Live Migration 6 of 6 – Completion



ldmd A destroys the domain on the source machine

# Live Migration Considerations 1/2

- Source Domain
  - Can not own physical devices, must use virtual I/O
  - Power Management in “performance” mode (the default)
  - Must be running LDomS 2.1 & updated firmware
    - (i.e. 7.4 or 8.1)
  - Must be running S10U9 or later
- Target System
  - must have enough resources (cpu, mem)
  - must have appropriate VIO services (vds, vsw, vcc)
  - must be able to provide required VIO devices (vdisk, vnet)
  - must be cpu-compatible:
    - Same processor type, same clock frequency (for now)

## Live Migration Considerations 2/2

- Secure memory transfer
  - while guest is operational
    - Transfer time depends on guest activity and memory content
    - Timings similar to warm migration
  - encrypted
    - provide at least one crypto unit to Control Domain
- Typical guest suspend time < 1s
  - Depends on guest activity
  - Requires about 1 core (T3)
- Live Migration not a replacement for HA
- White paper on best practices to be published soon

# Agenda

- Very Short Introduction to LDoms
- What's New in Version 2.1
  - Secure Live Migration
  - **Dynamic Reconfiguration & Power Management**
  - Inter-Vnet LDC Channels
- Management with OpsCenter
- LDoms in SuperCluster
- On or Behind the Horizon...




# Dynamic Resource Allocation

Completely Dynamic since Version 2.0

| Feature       | Supported |
|---------------|-----------|
| CPU DR        | ✓         |
| Memory DR     | ✓         |
| Virtual IO DR | ✓         |
| Crypto DR     | ✓         |

# Automatic Dynamic Resource Management





# SPARC T3 Power Management Enhancements

- CPU Clock Speed Adjustments
  - Increase or decrease clock speed based on CPU utilization
- Memory Power Management
  - Put under-utilized memory in a deeper idle mode
- Power Limit
  - Set a power limit for the system
  - Reduced the power state of manageable resource if the limit is reached

**Benefit:** Better controls on power consumption

# Agenda


- Very Short Introduction to LDoms
- What's New in Version 2.1
  - Secure Live Migration
  - Dynamic Reconfiguration & Power Management
  - **Inter-Vnet LDC Channels**
- Management with OpsCenter
- LDoms in SuperCluster
- On or Behind the Horizon...



# Inter-Vnet LDC Channels

Reduce LDC usage for complex network setups

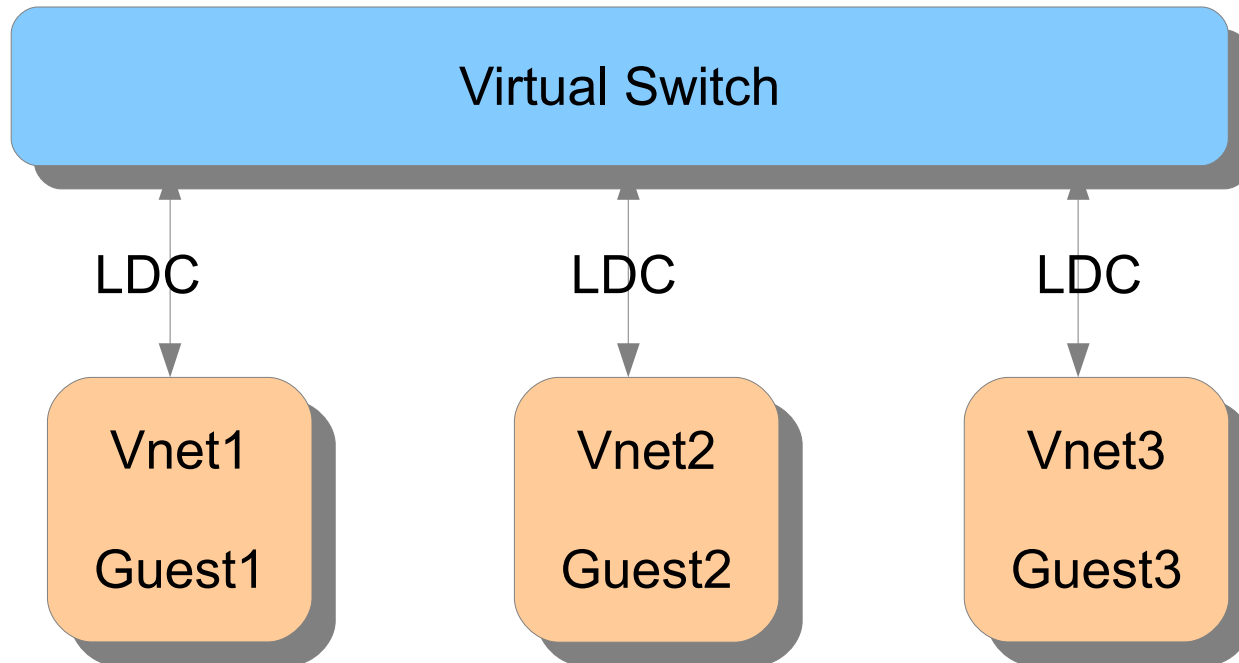
- Default behaviour: NxM LDCs



# Inter-Vnet LDC Channels

Reduce LDC usage for complex network setups

- New (optional) behaviour: LDC tree



# Inter-Vnet LDC Channels Details

- A new CLI option 'inter-vnet-link' is introduced to enable/disable the allocation of Inter Vnet LDC channels.
  - By default it is enabled(or ON)
  - This option is a Virtual Switch wide setting, that is enabling/disabling affects all Vnets in a given Virtual Switch.
  - Can be dynamically enabled/disabled without stopping the Guest domains.
  - The Guest domains dynamically handle this change.

## Other New Features in Release 2.1

- CPU Affinity
  - keep CPUs assigned to a guest on a minimal number of cores
- Virtual Device Service Validation
  - Reduce risk for misconfiguration
- Updated LDoms MIB
  - MIB now supports all 2.1 features (monitoring)
- Automatic Crypto-Removal during DR
  - Remove Crypto along with CPU
- P2V Enhancements


# Agenda

- Very Short Introduction to LDoms
- What's New in Version 2.1
  - Secure Live Migration
  - Dynamic Reconfiguration & Power Management
  - Inter-Vnet LDC Channels
- **Management with OpsCenter**
- LDoms in SuperCluster
- On or Behind the Horizon...



# Comprehensive Full-Stack Management

Applications to Disk





# LDoms & OpsCenter

| Feature                        | Support               |
|--------------------------------|-----------------------|
| Deploy on bare metal           | ✓                     |
| Manage what's been deployed    | ✓                     |
| Discover existing guests       | ✓ <b>Next Release</b> |
| Manage existing guests         | ✓ <b>Next Release</b> |
| DR on existing guests          | ✓ <b>Next Release</b> |
| Vpool on existing guests       | ✗ <b>Planned</b>      |
| Add storage on existing guests | ✗ <b>Planned</b>      |
| LM for existing guests         | ✗ <b>Planned</b>      |
| Redundant IO Domains           | ✗ <b>Planned</b>      |
| Automatic Failover             | ✗ <b>Planned</b>      |

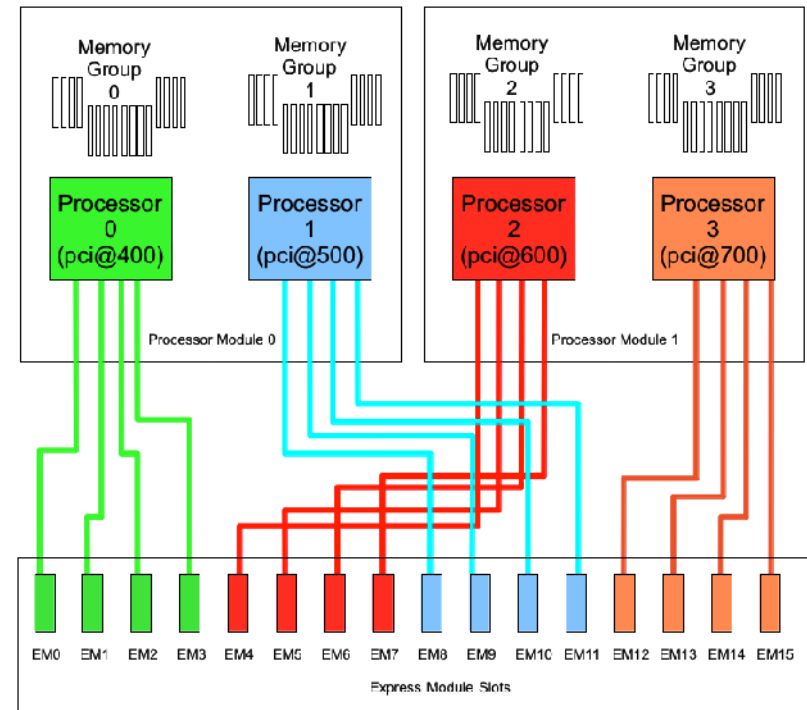
# Agenda

- Very Short Introduction to LDoms
- What's New in Version 2.1
  - Secure Live Migration
  - Dynamic Reconfiguration & Power Management
  - Other Improvements
- Management with OpsCenter
- **LDoms in SuperCluster**
- On or Behind the Horizon...



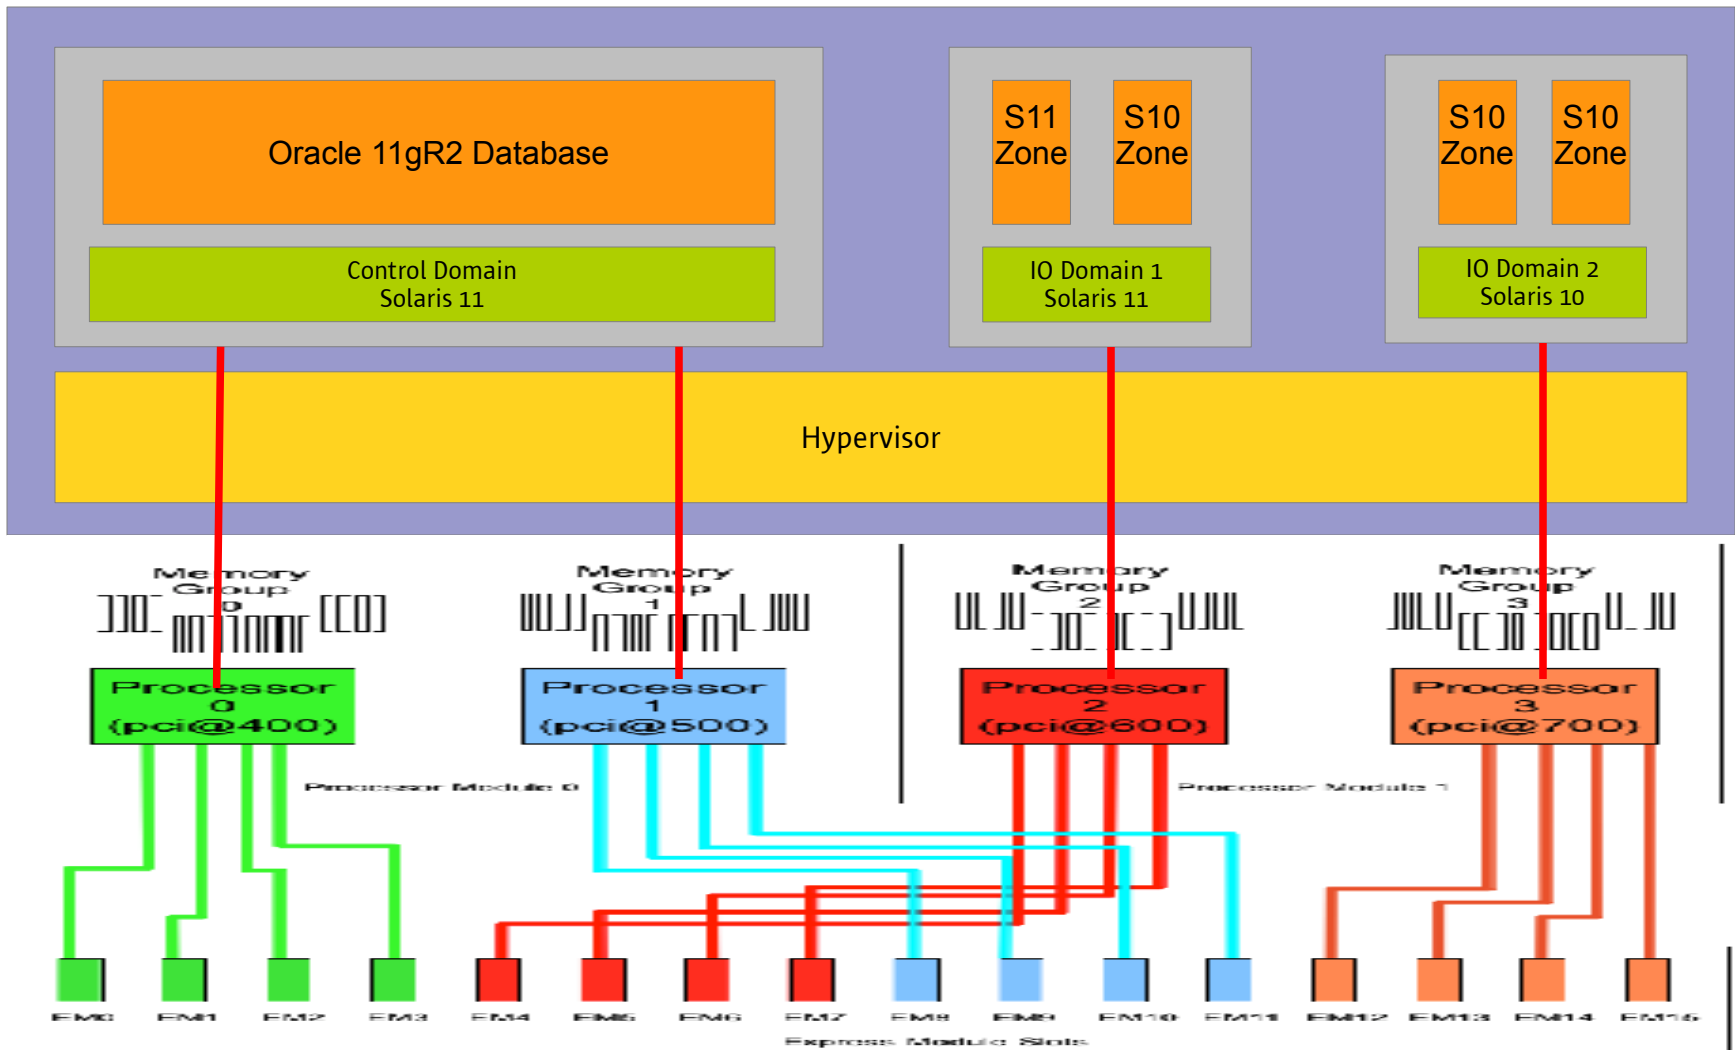
# LDoms for “Hardware Partitioning” in SuperCluster

- T4-4:
  - 4 CPUs
  - 1 PCIe Root Complex/CPU
  - 1 Memory “Group”/CPU
- SuperCluster
  - Partitioned along Root Complex Boundaries
  - IO “fixed” per domain
  - Implemented with “ldm add-io”
  - CPU/Memory more flexible



# LDoms for “Hardware Partitioning” in SuperCluster

## A typical Example



# Future Enhancements

Oracle does not allow us to talk about future releases, but....

- ...there are some things we are thinking about:
  - Continue to support future SPARC systems
  - I/O Virtualization – SR-IOV
  - Management via Oracle VM Manager
  - Improvements for redundant IO Domains
    - Multipathing failback & management
    - Support by OpsCenter
  - Fewer restrictions for Live Migration
    - Different CPU types
    - whole core constraint
  - etc...

# OVM Server for SPARC Key Features

- Secure Live Migration
- Dynamic Reconfiguration
- Redundant Virtual IO
  - Redundant IO Domains
- Static Direct IO
- Virtualized Crypto Acceleration
- Supported with
  - RAC
  - Solaris Cluster
- Lowest Overhead
- No Additional Cost

# For more Information

- <http://www.oracle.com/virtualization>
- <http://www.oracle.com/sun>
- Oracle VM OTN Forum
- Oracle Virtualization Blog, more...

The image shows a screenshot of the Oracle's Virtualization website. The main heading is "Oracle + Sun Virtualization Strategy". Below it, there is a section for "Oracle Virtualization" with a "Learn more >>" link. The page lists various virtualization solutions, including Oracle VM Server for x86, Oracle VM Server for SPARC, Oracle Solaris Containers, Desktop Virtualization, and Storage Virtualization. To the right of the website screenshot is a diagram illustrating the Oracle VM Server for SPARC architecture. The diagram shows a "sun2v virtual machine" containing three "Java App" icons and a "Solaris" operating system icon. This virtual machine is managed by a "Hypervisor" layer, which sits on top of the "SPARC CMT hardware".



Stefan.Hinker@oracle.com  
<https://blogs.oracle.com/cmt>



# **Hardware and Software Engineered to Work Together**

**ORACLE®**