
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.

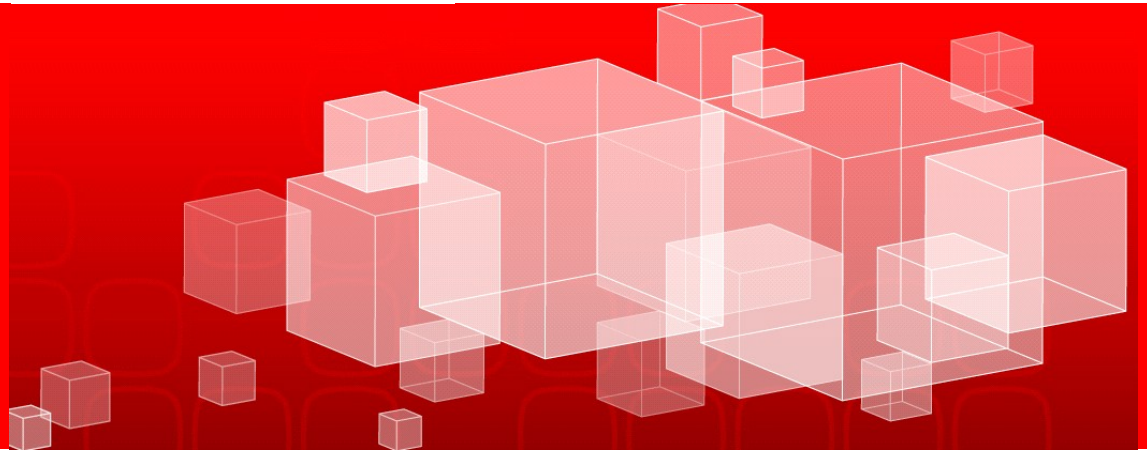

ORACLE®

Virtualisierte Netze mit Oracle VM Server for SPARC Guest Domains

Hartmut Streppel
Principal Sales Consultant, Server Architects Northern Europe

Agenda

- Oracle VM Server for SPARC
- Virtuelle Netze in OVM for SPARC
  - Konfiguration in einer Service Domain
  - Konfiguration in einer Guest Domain
- Zusammenfassung

Oracle VM Server for SPARC

Best Virtualization for Your Enterprise Workloads

Optimized for SPARC & Oracle Solaris

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

Firmware-based hypervisor

Each logical domain runs in dedicated CPU thread(s) on dedicated memory

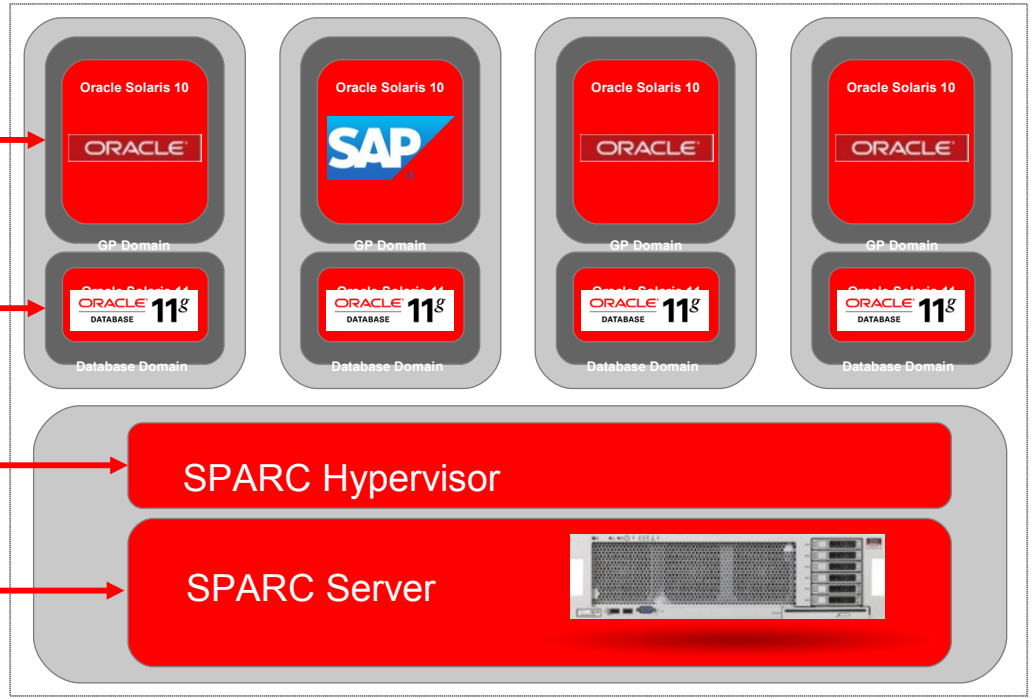

Logical Domains

- SPARC Hypervisor

- Immer da!
- Verwaltet Rechte an der Hardware
- Kein Overhead
- “primary” Domain besitzt initial alle Hardware

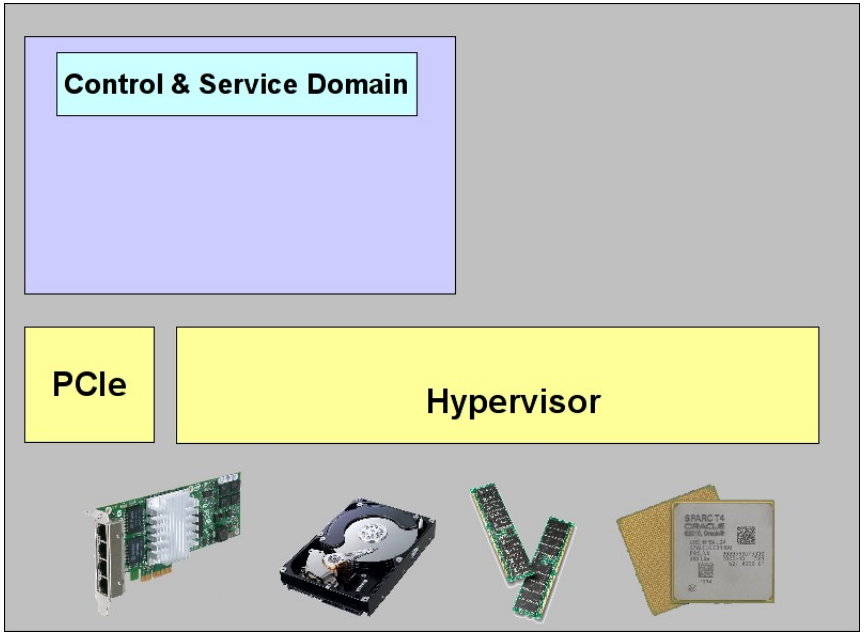

Domain Components

SPARC Hypervisor

- Gast LDOMs habe immer(!) eigene CPUs und Memory
- können eigene IO Geräte besitzen
- können virtuelle IO Geräte nutzen
- Virtuelle Geräte werden von IO Domains zur Verfügung gestellt

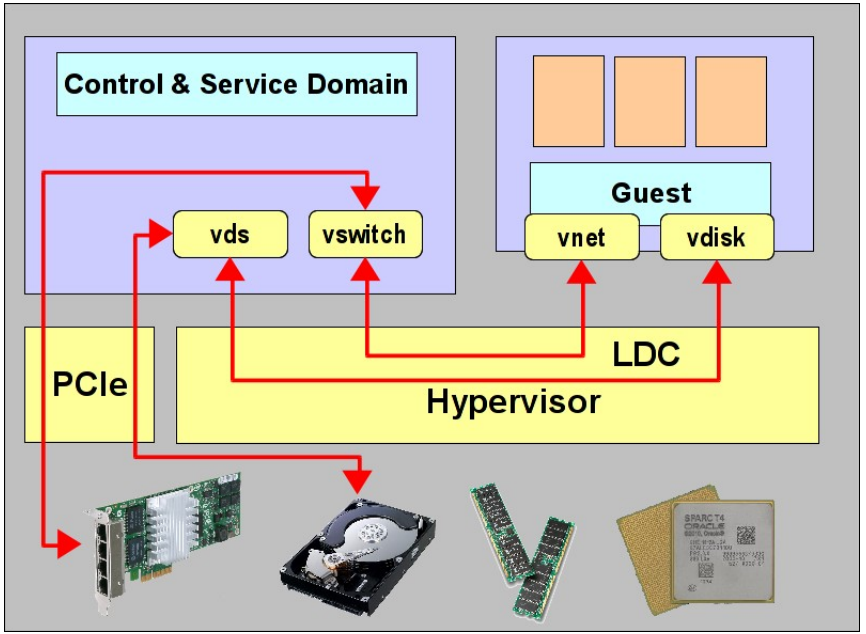

Konfiguration virtueller Netze in Solaris 10

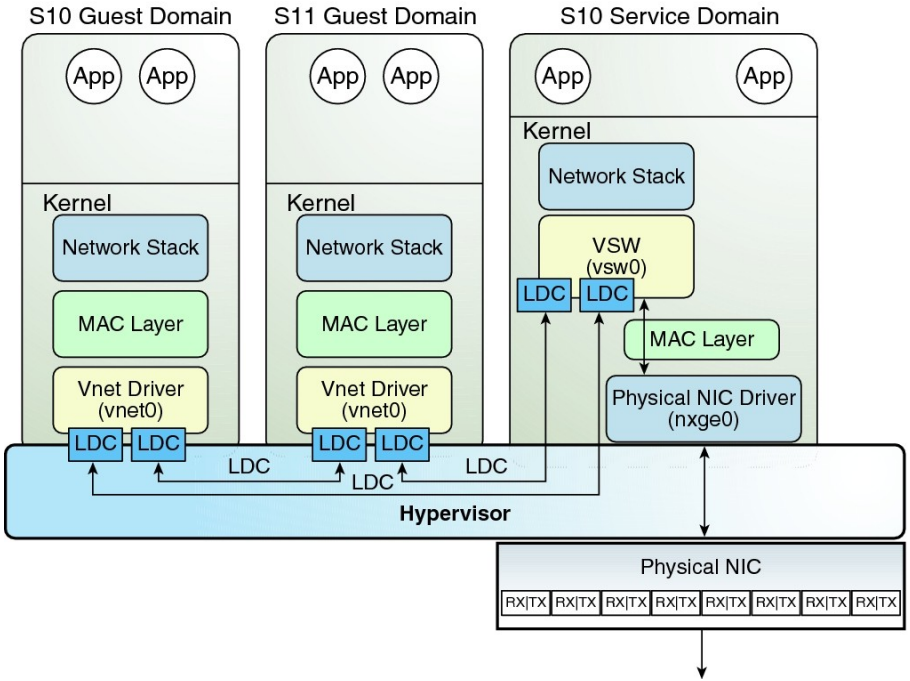

Konfiguration virtueller Netze in Solaris 11

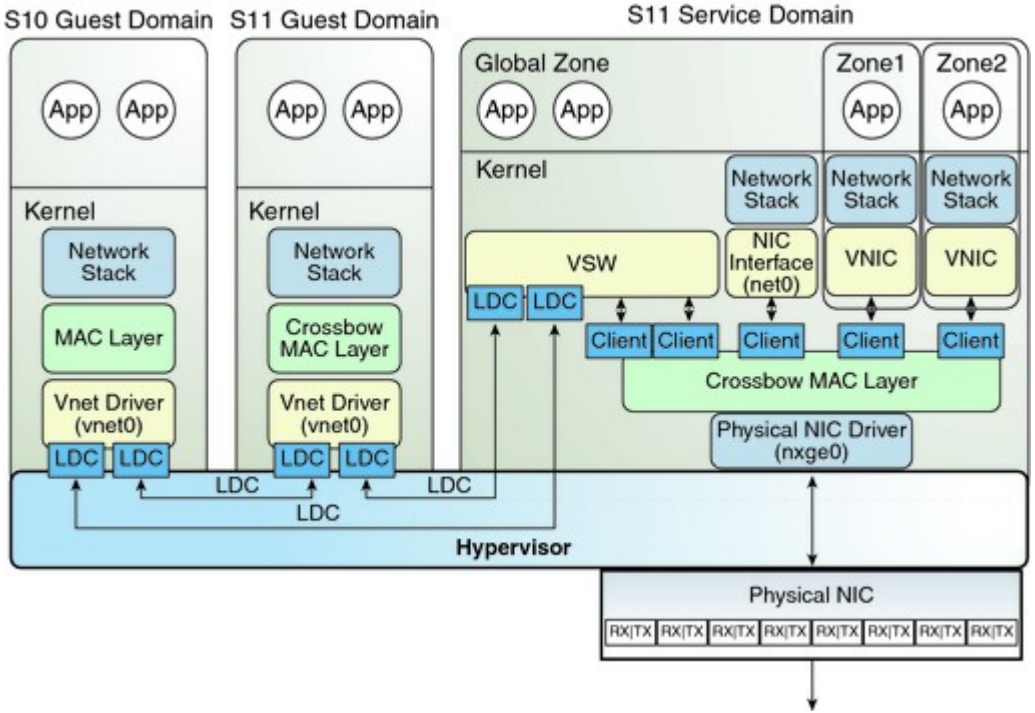

Logical Domain Channels - LDC

LDC

- LDCs stellen die Verbindung her zwischen Services in der Gast und einer anderen Domain
- Anzahl der LDCs pro Domain ist begrenzt

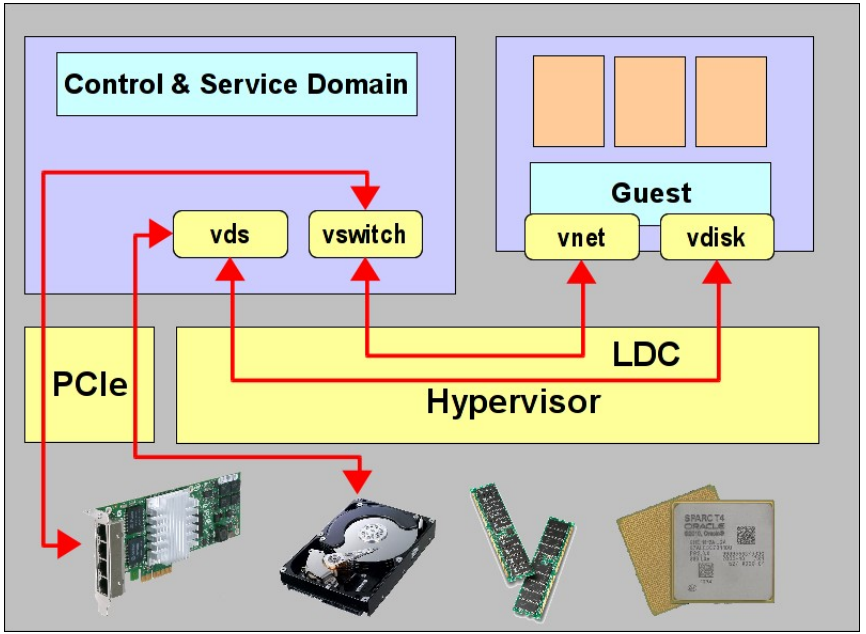

Logical Domain Channels

- Limitierung der LDCs durch Parametrierung der vswitches
  - **`ldm add-vswitch … inter-vnet-link=off …`**
  - **`ldm set-vswitch … inter-vnet-link=off …`**
- `=off`: Performance-Einbuße möglich

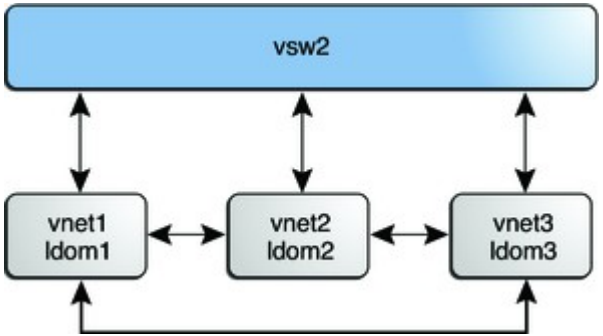

Vswitch mit Inter-vnet-links

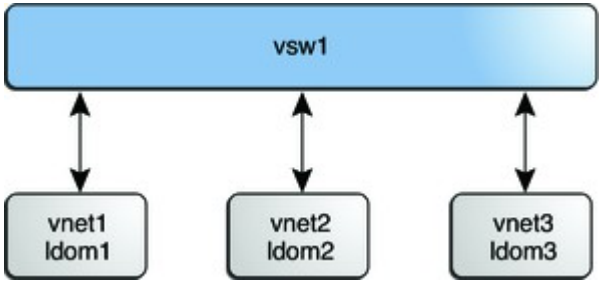

Vswitch ohne Inter-vnet-links

What's New in LDoms 3.1

- SR-IOV for non-primary Root Domains
- SR-IOV for Infiniband
- Dynamic SR-IOV
- Virtual networking performance improvements (now wire speed…)
- X-CPU Migration support for T5, M5 and M10
- New command: ldmpower (monitor power usage per domain)
- Auto-Recovery

Agenda

- Oracle VM Server for SPARC
- Virtuelle Netze in OVM for SPARC
  - In einer Service Domain
  - In einer Guest Domain
- Zusammenfassung

Konfiguration virtueller Netze

- Konfiguration in einer oder mehreren Service Domains
  - Physikalischen Adapter (NIC) (optional)
  - Virtuelle Switches (vswitch)
  - Virtuelle Adapter (vnet)

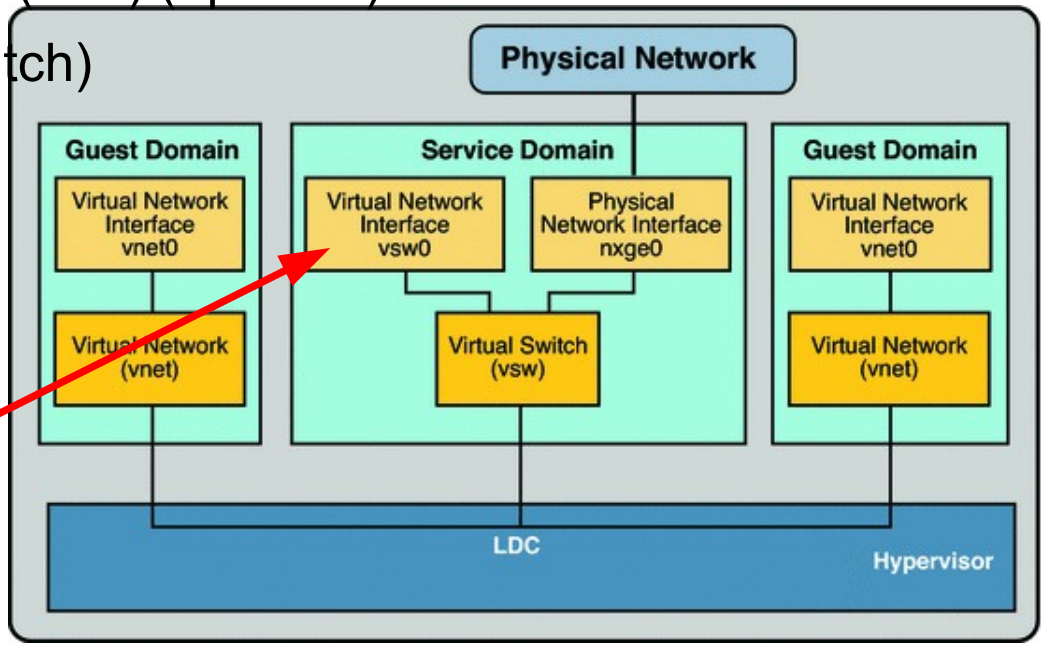

Achtung: vsw0 muss unter Solaris 10 explizit angelegt werden

Konfiguration

- Anlegen des vswitch

```
# ldm add-vsw net-dev=net0 vsw0 primary
```

- Anlegen des vnet

```
# ldm add-vnet linkprop=phys-state vnet0 vsw0 guest1
```

- Weiterleitung des Link-Status
  - linkprop=phys-state
  - Nicht default!

Konfiguration mehrerer virtueller Netze

- Anwendungsfall: 2 physikalische Cluster Interconnects
- Anlegen der zusätzlichen vswitch'e

```
# ldm add-vsw mode=sc net-dev=net1 priv1 primary
# ldm add-vsw mode=sc net-dev=net2 priv2 primary
```

- Anlegen der vnet's

```
# ldm add-vnet linkprop=phys-state vnet2 priv1 cluster1
# ldm add-vnet linkprop=phys-state vnet3 priv2 cluster1
```

- `mode=sc` - Cluster Interconnect benötigt höchste Priorität für Heartbeats

Konfiguration mit „device number“

- Spezifikation der „device number

```
# ldm add-vnet id=7 vnet7 vsw7 guest1
```

- Überprüfung in der Service und der Gast Domain

```
# ldm list-bindings primary
NAME             MAC                 NET-DEV   ID   DEVICE
         vsw7  00:14:4f:fb:b7:02                3    switch@3
           PEER                          MAC
        vnet7@guest1              00:14:4f:f9:47:08
# ldm list-bindings guest1
 NAME    SERVICE                  ID   DEVICE     MAC
    vnet7  vsw7@primary           7    network@7  00:14:4f:f9:47:08
```

Agenda

- Oracle VM Server for SPARC
- Virtuelle Netze in OVM for SPARC
  - In einer Service Domain
  - In einer Guest Domain
- Zusammenfassung

Konfiguration eines NICs in Gast Domain

- Keine Änderung zur IO- oder Root-Domain, außer
  - `dladm show-phys` zeigt als DEVICE „vnet“

```
[guest]# dladm show-phys
LINK    MEDIA     STATE     SPEED     DUPLEX    DEVICE
net0    Ethernet  up          0       unknown   vnet0
net1    Ethernet  up          0       unknown   vnet1
```

VNICs auf VNETs

- Konfiguration von S11 vnic's auf OVM vnet's

```
[gastdomain] # zoneadm -z vm10 boot
zone 'vm10': failed to create vnic for net0: operation failed
zoneadm: zone 'vm10': call to zoneadmd failed
```

- Nicht möglich bis OVM 3.0.0.2
  - Grund: fehlende MAC-Adressen für VNICs
- https://blogs.oracle.com/jsavit/entry/vnics_on_vnets_now_available
- Anlegen zusätzlicher MAC-Adressen für den vswitch

VNICs auf VNETs

- Setzen der MAC-Adressen, z.B.

```
# ldm add-vnet vnet1 primary-vsw0 mydomain1
# ldm set-vnet alt-mac-addrs=auto,auto,auto,auto vnet1 mydomain1
```

- Dedizierte MAC-Adressen möglich:
  - 00:14:4F:F8:00:00 ~ 00:14:4F:FF:FF:FF (allowed for LDOMs)
  - 00:14:4F:F8:00:00 – 00:14:4F:FB:FF:FF (automatic allocation)
  - 00:14:4F:FC:00:00 – 00:14:4F:FF:FF:FF (manual allocation)
  - https://docs.oracle.com/cd/E35434_01/html/E23807/assignmacaddressesautomaticallyormanually.html#scrolltoc
- Duplicate MAC Address Detection

IPMP

- „IPMP“ in einer Service Domain über zwei NICs
  - Nicht empfohlen; komplexes Setup
  - Link Aggregation verwenden für vswitch
- IPMP in der Gast Domain
  - Über zwei NICs von einer Service Domain
  - Über zwei NICs von zwei Service Domains (beste Verfügbarkeit)

Link Aggregation in der Service Domain

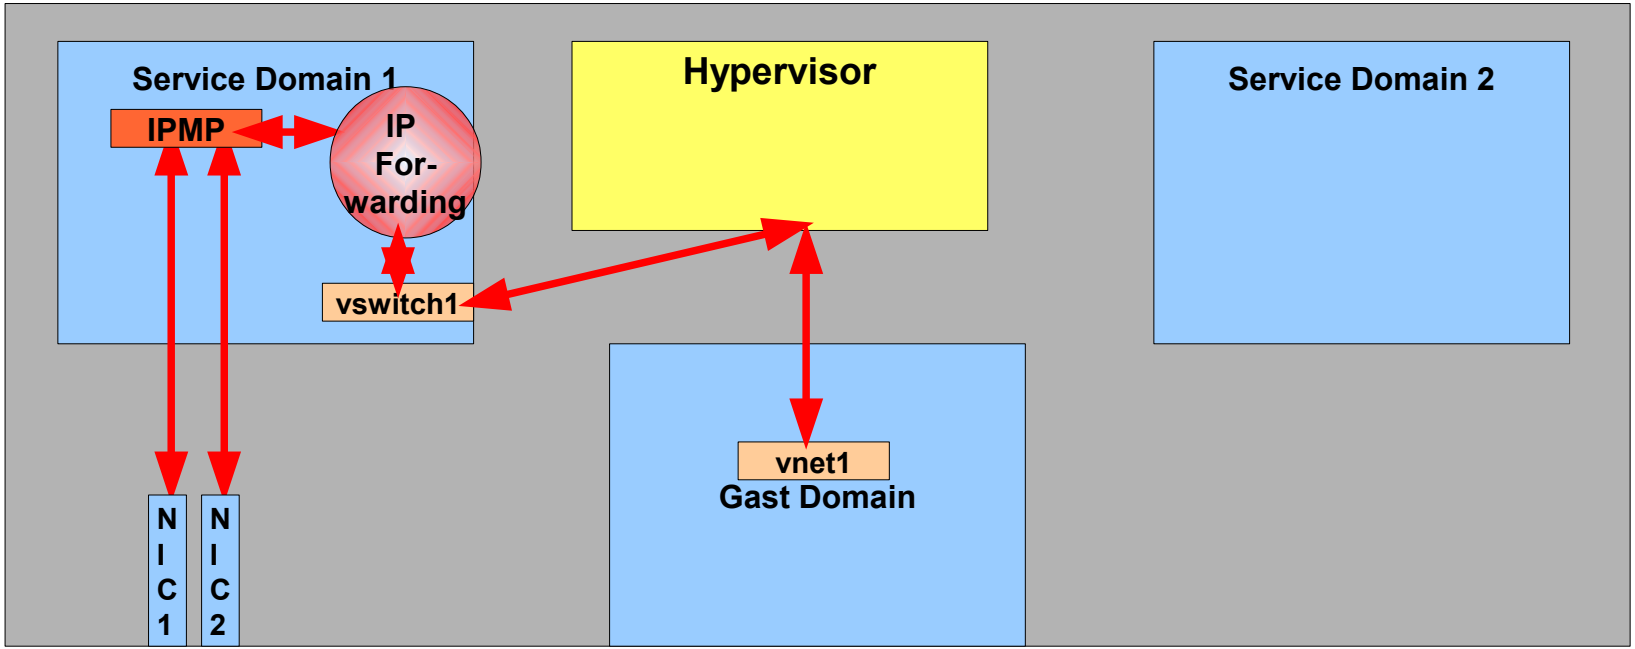

Link Aggregation in der Service Domain

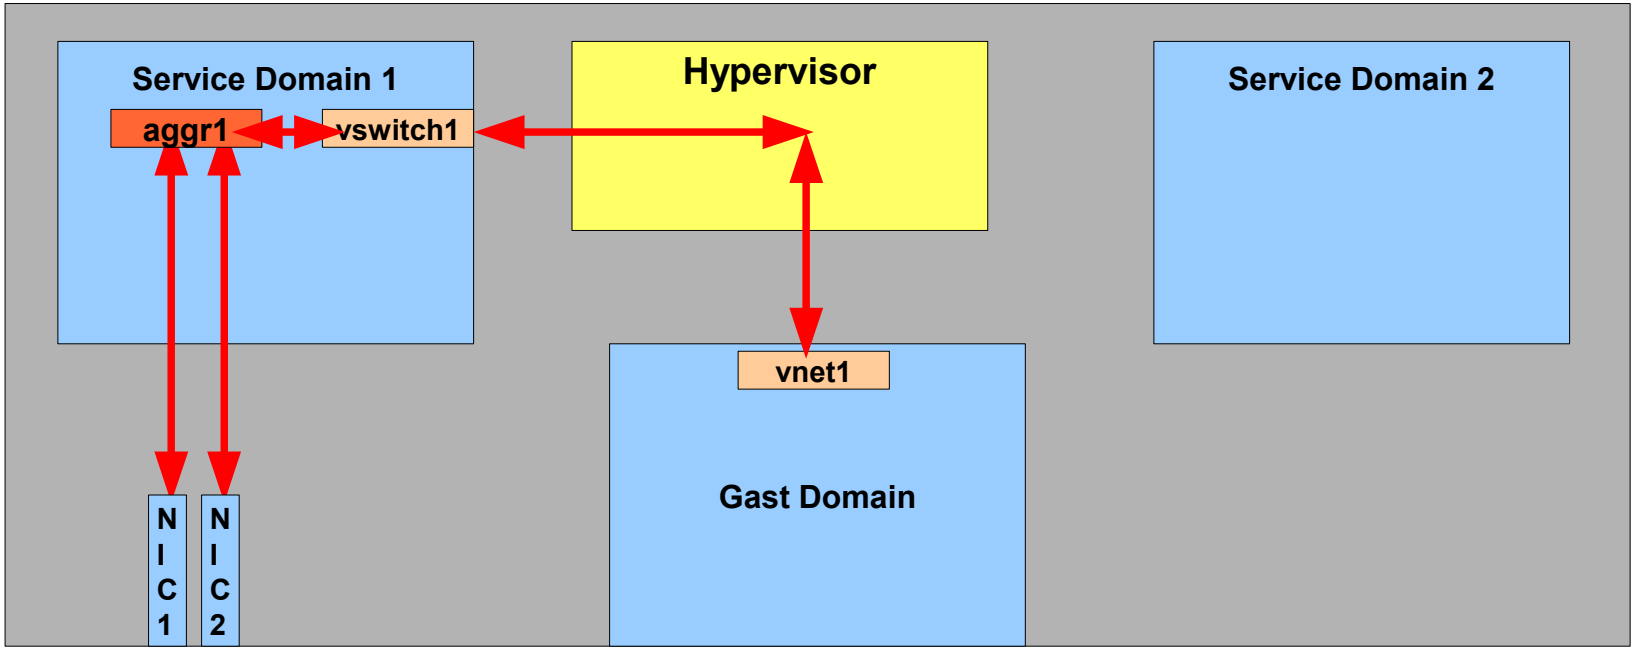

IPMP in der Gast Domain

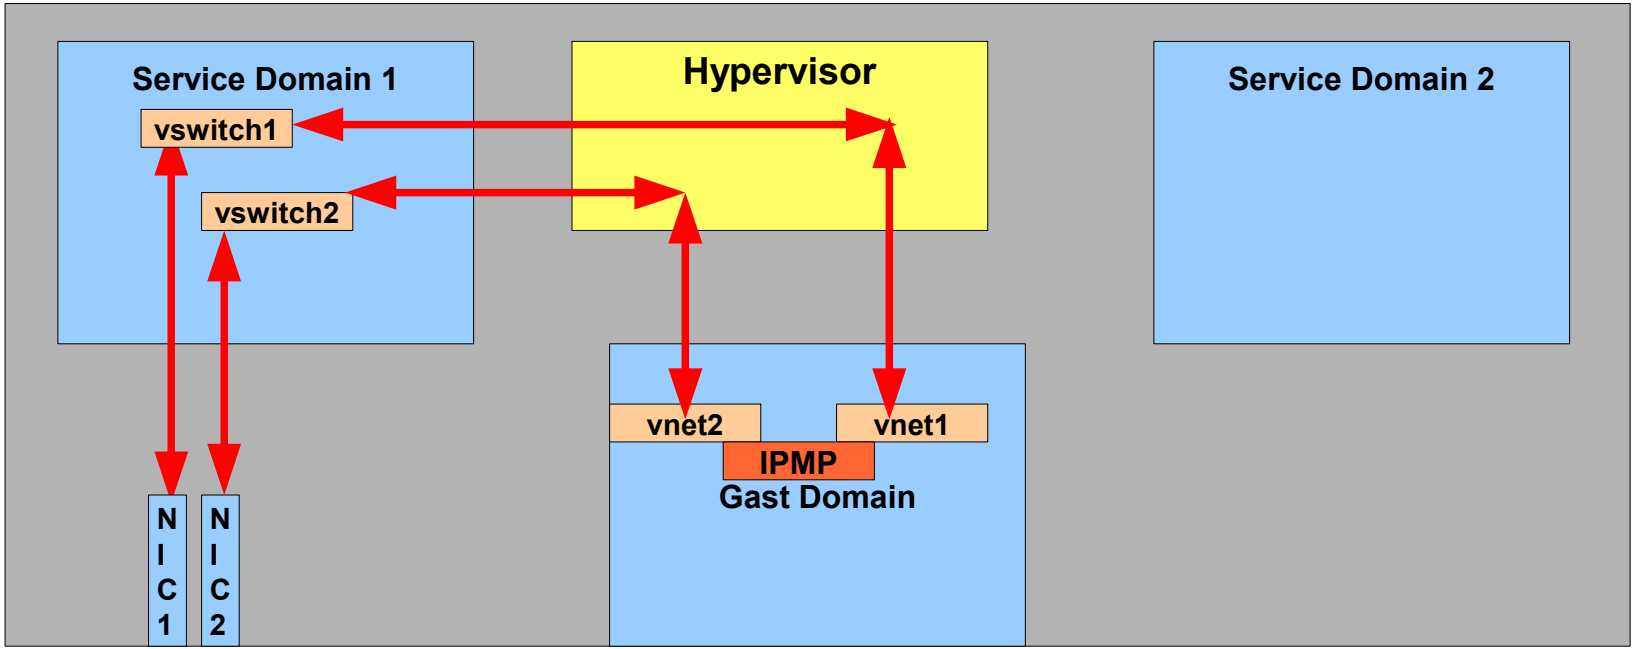

IPMP in der Gast Domain

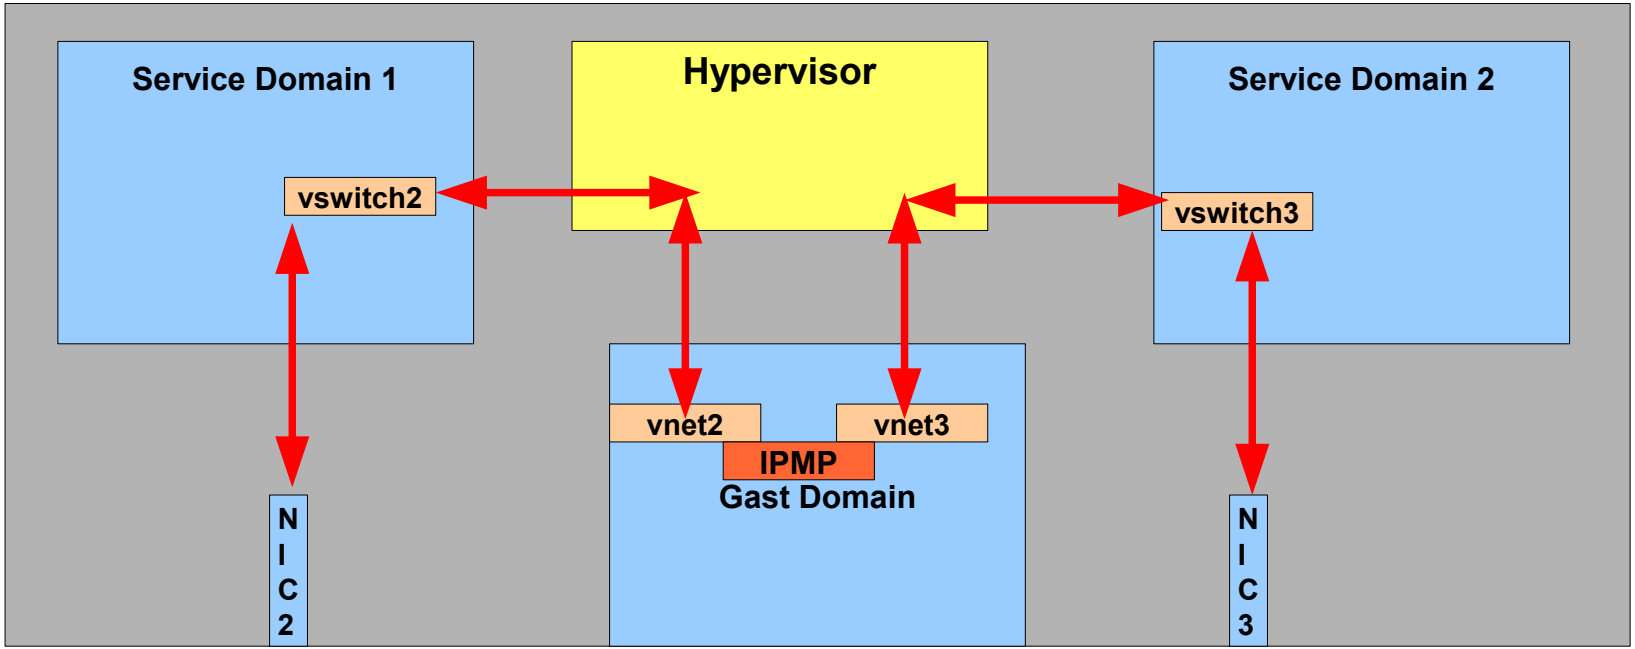

Network Performance in Gast Domains

- Nutzen Sie OVM 3.1
- extended_mapin_space=on (default ab 3.1)
- https://blogs.oracle.com/raghuram/entry/ldoms_virtual_network_performance_greatly1

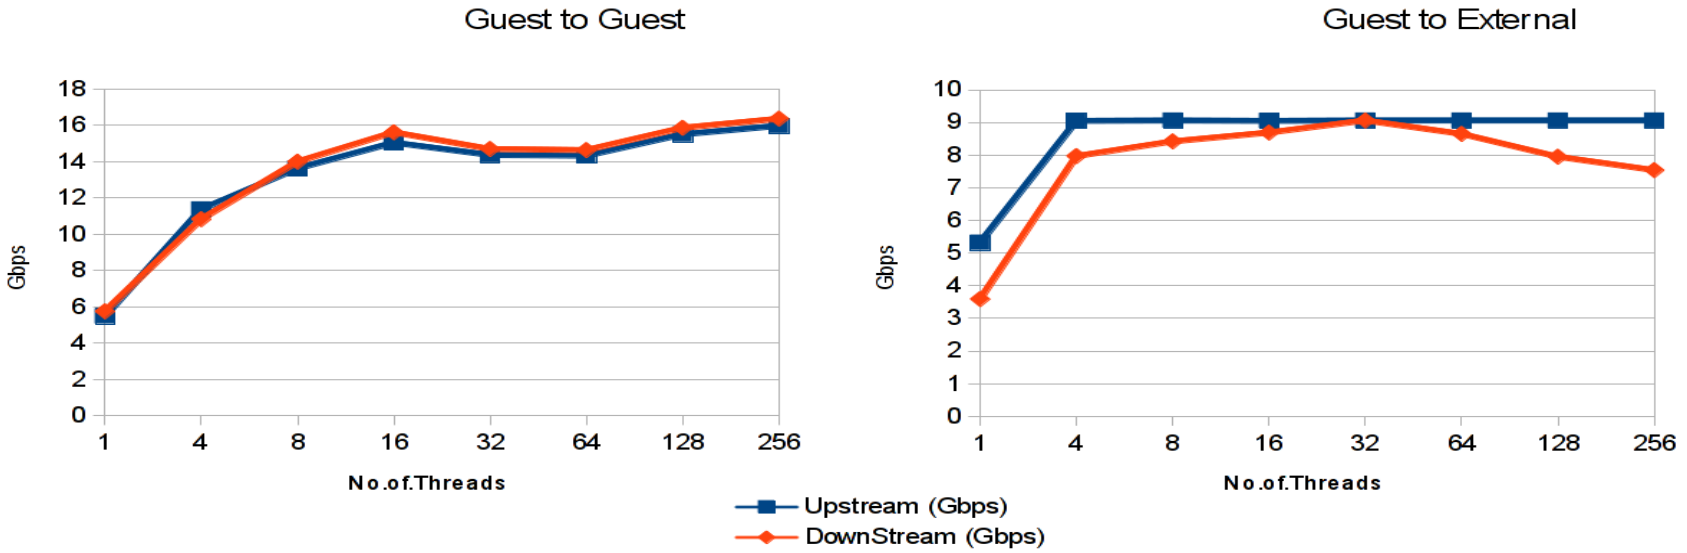

Live Migration Support

- „Each virtual network device in the domain to be migrated must have a corresponding virtual network switch on the target machine.“
- „Identische“ Netz-Infrastruktur notwendig in der Zielumgebung

Agenda

- Oracle VM Server for SPARC
- Virtuelle Netze in OVM for SPARC
  - In einer Service Domain
  - In einer Guest Domain
- Zusammenfassung

Zusammenfassung

- Effiziente und performante Kommunikation zwischen Logical Domains
- Einfache Konfiguration
- Nahtlose Integration zwischen virtuellen Netzen in Solaris und virtuellen Netzen in Oracle VM for SPARC

Q&A

Hartmut.Streppel@oracle.com

ORACLE


ORACLE®