

Manage of OVM server Sparc with the VM manager, OpsCenter or VDCF

SOUG DAY 02.MÄRZ 2016

Roman Gächter



BASLE ■ BERN ■ BRUGG ■ DÜSSELDORF ■ FRANKFURT A.M. ■ FREIBURG I.BR. ■ GENEVA
HAMBURG ■ COPENHAGEN ■ LAUSANNE ■ MUNICH ■ STUTTGART ■ VIENNA ■ ZURICH

trivadis
makes IT easier. ■ ■ ■

■ Our company.

Trivadis is a **market leader in IT consulting, system integration, solution engineering** and the provision of **IT services** focusing on _____ and _____ technologies in Switzerland, Germany, Austria and Denmark. We offer our services in the following strategic business fields:



Trivadis Services takes over the interacting operation of your IT systems.

■ With over 600 specialists and IT experts in your region.

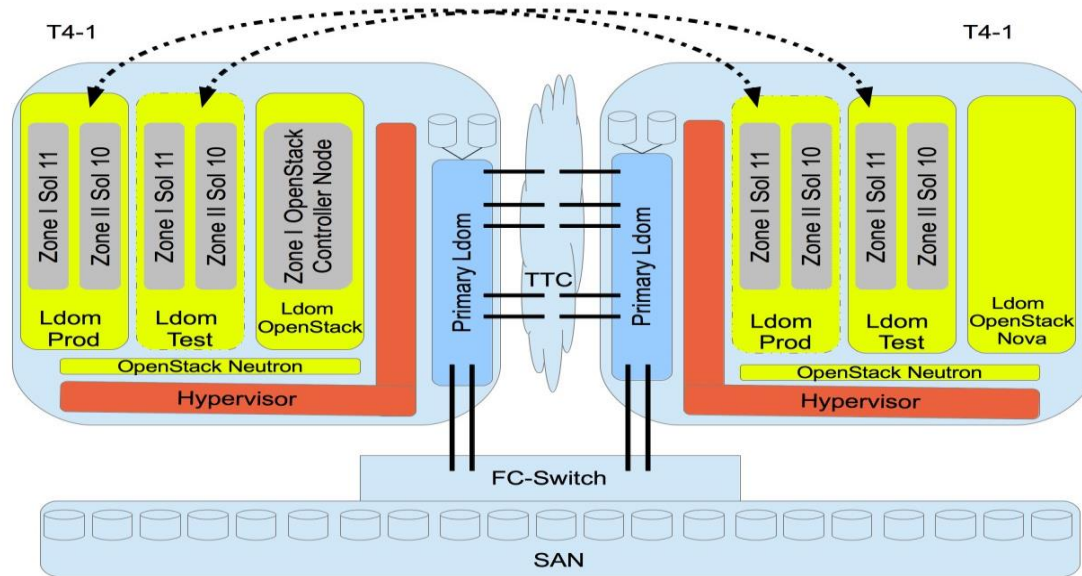


- 14 Trivadis branches and more than 600 employees
- 200 Service Level Agreements
- Over 4,000 training participants
- Research and development budget: CHF 5.0 million
- Financially self-supporting and sustainably profitable
- Experience from more than 1,900 projects per year at over 800 customers

Starting position

Starting position

New HW for test- and development-environment



■ Starting position

Criteria for rating I

- Support of our architecture
 - With current versions from Solaris and VM-servers
 - Support from current VM manager in combination with VM-Server X86
- Painless takeover from our consisting configuration
- HW monitoring from the SPARC T4 systems
- Configurationmanagement of the SPARC Hypervisors
- Configurationsmanagement from Solaris zones
- Optimal support of networkstacks in the hypervisor and in Solaris

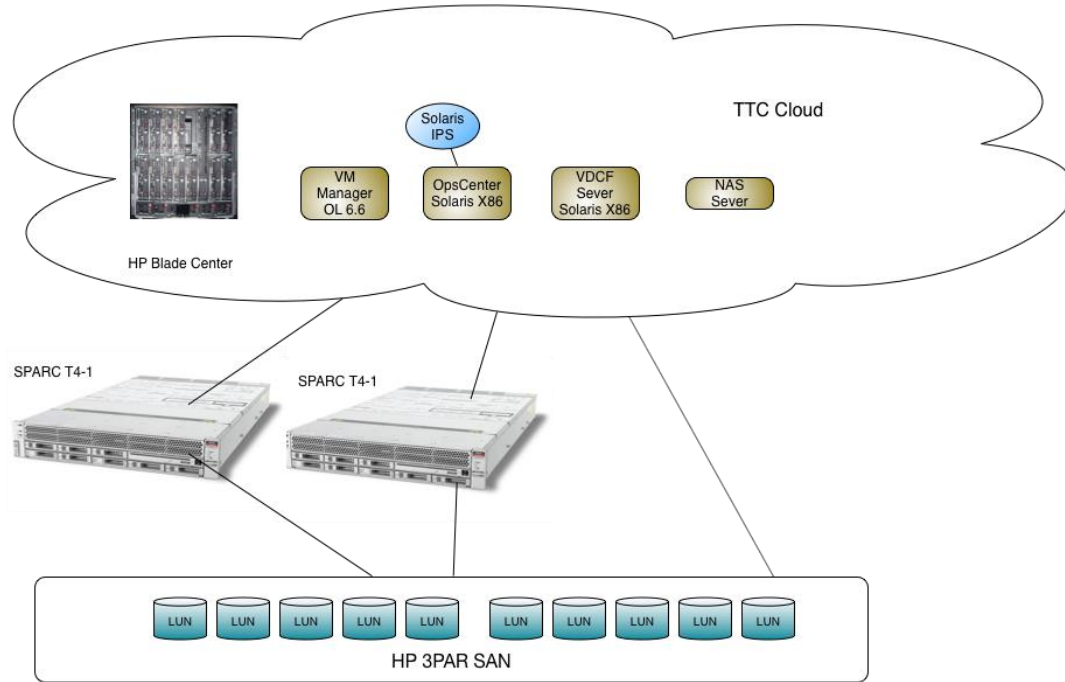
■ Starting position

Criteria for evaluation II

- Provisioning from Solaris and Solaris Zones
- Solaris updates / patching
- Support from our operational concepts
 - Migration from Zones / LDOMs of a HW to the other
- Monitoring LDOMs, zones and Solaris
- Simple Setup and operability of the tools
- The installation and the setup are not allowed to take too much time

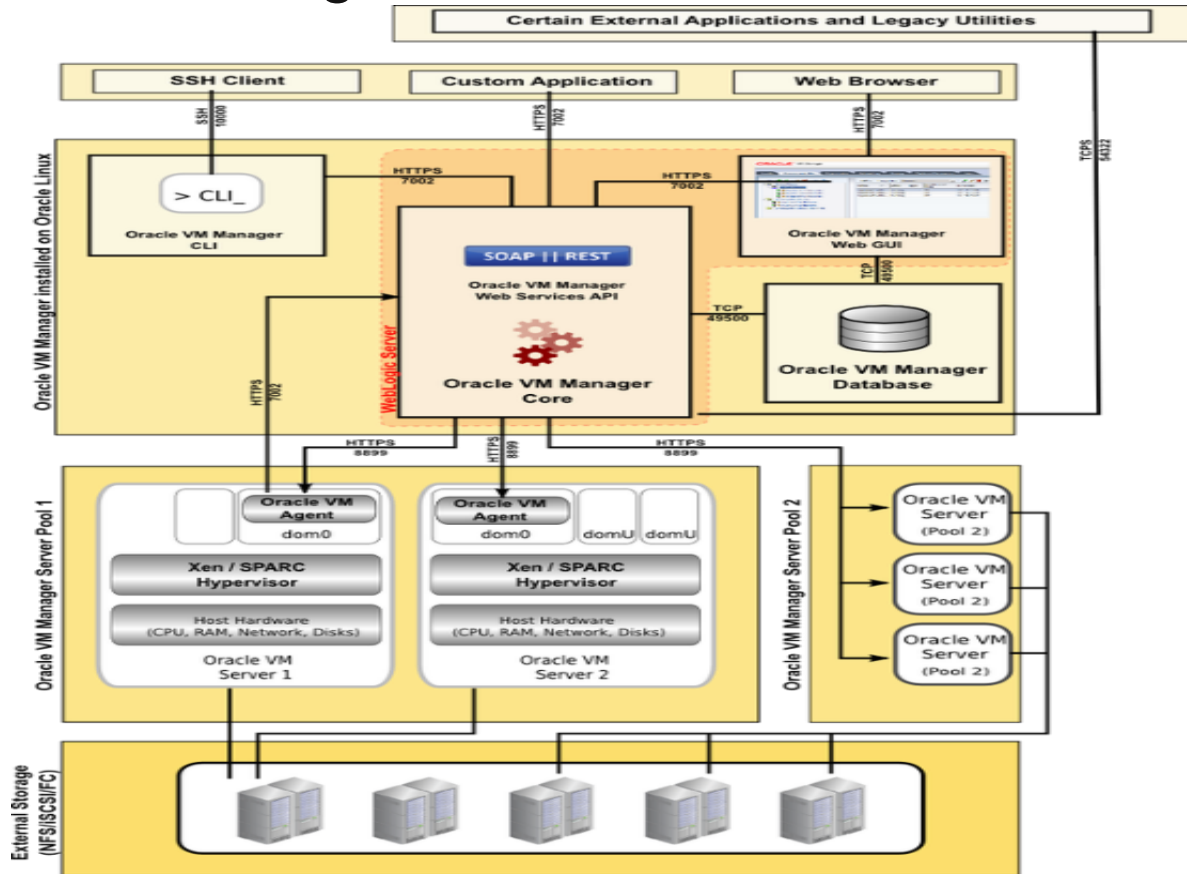
■ Trivadis technology center

Test scenario



Oracle VM manager

VM manager



■ VM manager

First impression

- Start with factory default LDOM configuration
- OVM server cluster
 - Repository and server pools only on NFS Share
 - OCFS2 only of Oracle Linux

■ VM Manager

Setup

- Controller domain
 - OVM server SPARC is a component of Solaris 11
 - Agent for VM manager necessary
 - Start with „factory default,, LDom configuration
 - Start with „DefaultFixed network configuration profile,,
- VM manager

```
# ./runInstaller.sh
```

■ VM Manager

Configuration

- Configuration sequence



VM manager

Configuration

- Example link
 - Aggregation
 - Trunk
 - Datalink multipathing

Edit Port

Server: germany
Port: bond2 on germany
Bonding: Link Aggregation
Addressing: None
IP Address:
Mask:
* MTU: 1500
Description:

Available Ports
net7 on germany (net7)

Selected Ports
net4 on germany (net4)
net5 on germany (net5)

Cancel OK

```
root@germany:~# dladm show-aggr bond2
LINK          MODE POLICY  ADDRPOLICY  LACACTIVITY LACPTIMER
bond2        trunk L4    auto        off         short
```

VM Manager

Configuration

- Network on hypervisor

Name	ID	Intra-Network Server	Network Channels					Description
			Server Management	Cluster Heartbeat	Live Migrate	Storage	Virtual Machine	
172.16.64.0	ac104000		√	√	√			
AGGR_1	109d7b930a						√	
TRUNC_PUBLIC	10e754566e						√	
VMIntraConnect	10db0999f9						√	InterConnect between T4
VMPublic	106b4d601d						√	

■ VM manager

Configuration

- Aggregate of «primary» and Vnic of «guest» LDOM

VSW

NAME	MAC	NET-DEV	ID	DEVICE	LINKPROP
DEFAULT-VLAN-ID	PVID VID	MTU	MODE	INTER-VNET-LINK	
109d7b930a	00:14:4f:fa:04:65	bond2	3	switch@3	

NETWORK

NAME	SERVICE	ID	DEVICE	MAC
MODE PVID VID	MTU	MAXBW	LINKPROP	
net1	109d7b930a@primary	1	network@1	00:21:f6:35:c7:d2

■ VM manager

Tested

- Repository and server pool on NFS Share
 - live migration is working
- Add disks (vdisk, local)
- Network
 - Bond ports
 - link aggregation
- Install Guest LDom with template
- Clone from LDom

■ VM manager

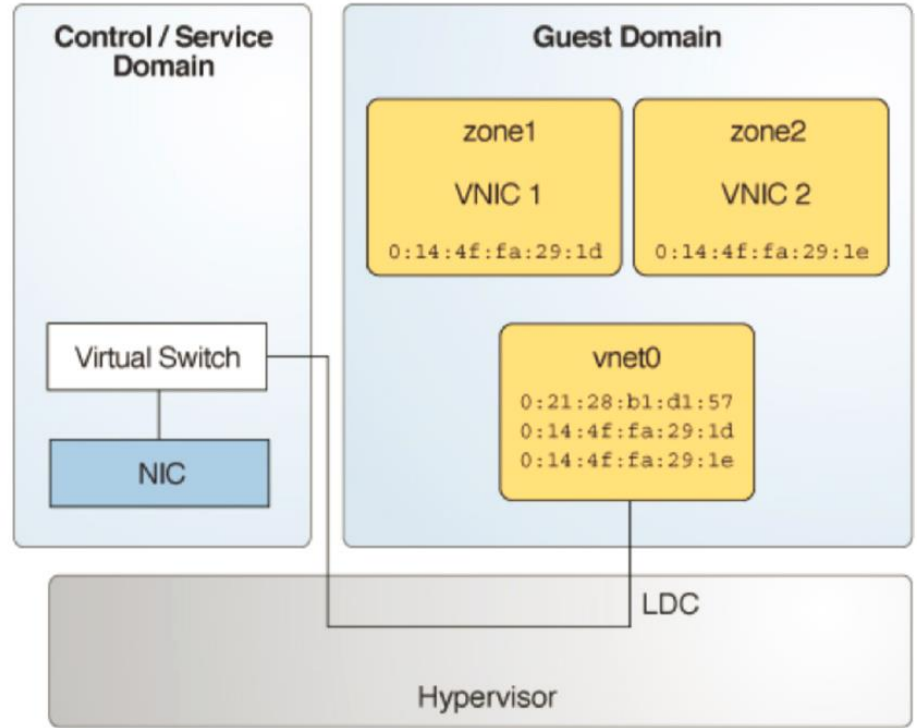
Remarks

- Solaris zones
 - He doesn't know
- SAN LUN's for cluster
 - Not supported
- Clone
 - Guest domain must be shut down

■ VM manager

Remarks

- Vnics over Vnet for zones
 - Like with the VM manager?



■ VM manager

Remarks

- No alternative mac`s about VM manager?
 - Attention, don't do it with ldm by hand!!!

(ldm set-vnet alt-mac-addr=auto,auto,auto,auto,auto vnet0 bern)

```
NETWORK
  NAME          SERVICE          ID  DEVICE  MAC
MODE  PVID  VID          MTU  MAXBW  LINKPROP
  vnet0        management-vsw0@primary  0   network@0  00:14:4f:fb:e0:6e
1
              1500          phys-state
                                00:14:4f:f8:22:5a
                                00:14:4f:f8:c3:c0
                                00:14:4f:fb:41:42
                                00:14:4f:f9:46:b3
                                00:14:4f:f8:6b:a8
```

■ VM manager

Remarks

- No manual changes to the LDoms configuration

```
rga@MacBook:~$ ssh -l root germany.ttc.trivadis.com
Last login: Sun Nov 15 19:28:31 2015 from 172.16.134.3
Oracle Corporation      SunOS 5.11      11.2      June 2014
```

This server is configured with the Oracle VM Agent. It has a specific configuration where LDoms virtual I/Os are disabled while Oracle VM Agent services have not been online.

To unconfigure the Oracle VM Agent and restore the regular behavior of LDoms virtual I/Os, refer to the Oracle VM Installation and Upgrade Guide, or read the `/README.OVM_UNINSTALL` file

Warning: making manual modifications in the management domain might cause inconsistencies between Oracle VM Manager and the server.

VM manager

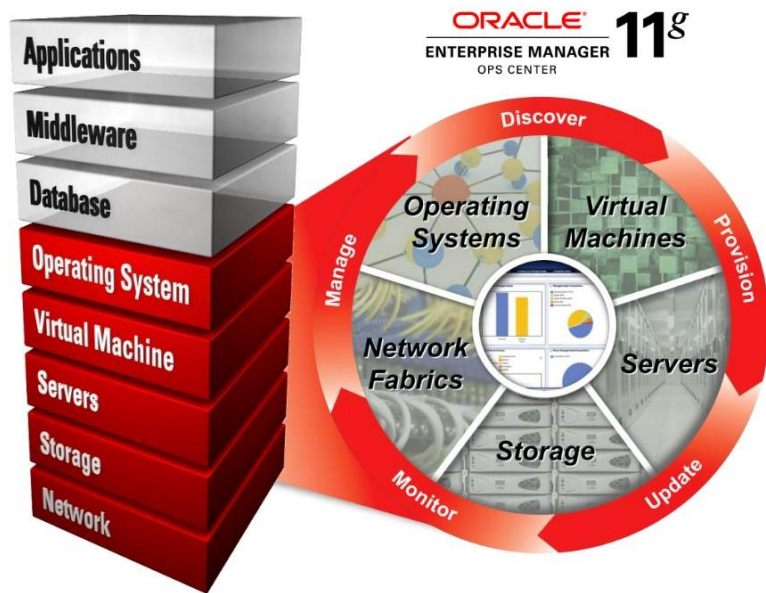
Results VM manager

Resultate von VM Manager		
Support unserer Architektur mit aktuellen Versionen von Solaris und OVM Server	Solaris 11.2 mit OVM Server 3.x sind unterstützt	+++
Schmerzfreie Übernahme unserer Bestehenden Konfiguration	LDOM und Netzwerkkonfigurationen gehen verloren	---
HW Monitoring der SPARC T4 Systeme	nein	---
Konfigurationsmanagement des SPARC Hypervisors	ja	++
Konfigurationsmanagement von Solaris Zonen	Zonen werden nicht erkannt	---
Optimale Unterstützung des Netzwerkstacks im Hypervisor und in Solaris	Linkaggregation ohne Modes, keine alternativ mac's	+
Provisioning von Solaris und Solaris Zonen	Über Templates?	
Solaris Updates / Patching	nein	---
Live migration von Zonen / Ldom's von einer HW auf die andere	Ja, mit NFS, Zonen nein	+
Monitoring von LDOM's, Zonen und Solaris	LDOM's ja, Zonen nein	+
Einfaches Setup und Bedienbarkeit der Tools	Relativ einfach	++

Oracle Enterprise Manager OpsCenter

■ OpsCenter

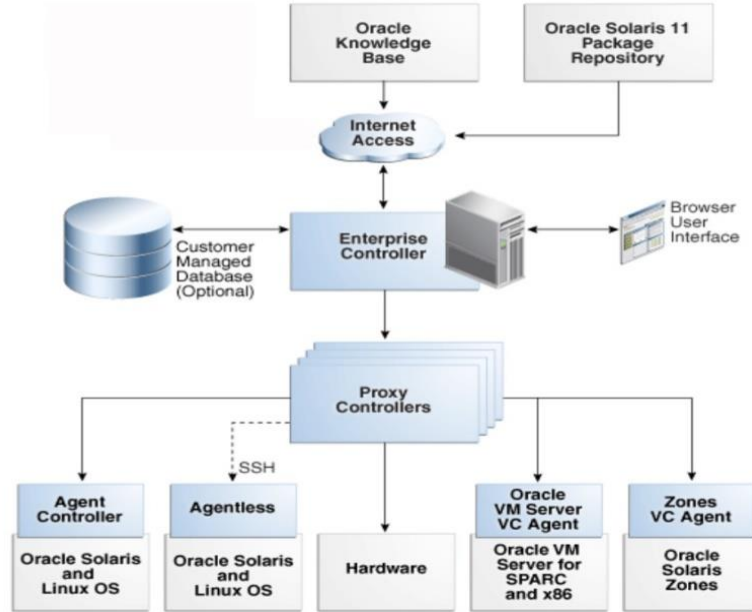
Overview



Integrated Infrastructure Management
+
Integrated Application-to-Disk Management
+
Integrated Lifecycle Management
+
Integrated Systems Management & Support

OpsCenter

Architecture



■ OpsCenter

First impression

- Current version?
 - OpsCenter 12c Release 3
 - OVM 3.3 not supported
 - Oracle VM Server SPARC 3.2
 - Oracle Virtual Machine X86 v3.0.x, 3.1.x and 3.2.x

■ OpsCenter

Setup

- Installations
 - Script for check the requirements
 - Script for installation
- **Setup**
 - «discovery» about existing VM manager
 - Additional agents
- **Error Messages**
 - Difficult to understand

■ OpsCenter

What went wrong?

Event Logs		Task Properties
Date ▲	Severity	Message
09/28/2015 1:01:05 pm UTC	WARNING	Model population failed for the ops01-e80b1b42/com.sun.hss.domain:name=172.16.65.37-global,type=ZoneServerContainer at 172.16.65.37 (172.16.65.37). ops01-e80b1b42/com.sun.hss.domain:name=172.16.65.37-global,type=ZoneServerContainer removed from the domain model. (10029)

OpsCenter

Awarded to few rights

The screenshot shows the 'Update Credentials' dialog box in Oracle Enterprise Manager Ops Center. The dialog has a title bar 'Oracle Enterprise Manager Ops Center - Update Credentials' and a sub-header 'Update Credentials' with a help icon and the Oracle logo. A note '* Indicates Required Field' is present. The form contains the following fields:

- Name:** orade
- Description:** Orade User
- SSH Section:**
 - Authentication Type:** Password (selected), Ops Center Key, Custom SSH Key
 - Login User:** oracle
 - Password:** (masked with dots)
 - Confirm Password:** (masked with dots)
 - Privileged Role:** root
 - Role Password:** (masked with dots)
 - Confirm Password:** (masked with dots)
 - SSH Port:** 22

Buttons for 'Update' and 'Cancel' are at the bottom right.

HW monitoring

Component Navigation

- System
 - Processors
 - Memory
 - Power
 - Cooling
 - Storage
 - Disk Controller
 - Disk Expander
 - Networking
 - PCI Devices
 - Service Processor

Summary Unassigned Incidents: ✖ 0 🔥 0 🔔 1

Model: SPARC T4-1

Server Name: switzerland

System Type: Rack Mount

Part Number: 30121114+1+1

Serial Number: 1146BDYC5E

System Identifier:

Management IP: 172.16.65.225

Management MAC Address: 00:21:28:D7:18:43

Power: On

Locator Light: Off

Actual Power Consumption: 318 watts

Operational Status: OK



HW monitoring

Subsystem Status		
Subsystem	Operational Status	Inventory
Processors	OK	Processors (Installed / Max): 1 / 1
Memory	OK	DIMMs (Installed / Max): 16 / 16
Power	OK	PSUs (Installed / Max): 2 / 2
Cooling	OK	Chassis Fans (Installed / Max): 12 / 12 PSU Fans (Installed / Max): - / -
Storage	OK	Internal Disks (Installed / Max): 4 / 8
Networking	OK	Installed Ethernet NICs: 4

Displaying 1 - 6 of 6

Firmware		
Description	Type	Version
Sun System Firmware	System-Firmware	8.6.0.b
Sun(TM) Integrated Lights Out Manager(Build 93935)	SP-Firmware	3.2.4.1
OBP firmware	OBP	4.36.2
Hypervisor firmware	Hypervisor	1.13.2.b

Displaying 1 - 4 of 4

HW monitoring



Monitoring

- A lot policies available

Summary	
Name:	Copy of OC - Global Zone
Description:	Generic Solaris Container
Target Type:	Operating Systems
Alert Monitoring Rules	Enabled ?
Appliance Health	Yes
Storage Library Usage Percentage	Yes
NAS Library Status	Yes
Number of Boot Environments	Yes
CPU Usage Percentage	Yes
Memory Usage Percentage	Yes
Swap Usage Percentage	Yes
File System Used Space Percentage	Yes
File System Reachability	Yes
System Load 15-Minute Average Per CPU Core	Yes
Disk IO Utilization Percentage	Yes
Disk IO Queue Length	Yes
Network Bandwidth Utilization Percentage	Yes
Process CPU Utilization Percentage	Yes
Operating System Reachability	Yes
ZPool Usage Percentage	Yes
Boot Env Usage Percent in a ZPool	Yes
SMF Service State	Yes
Service Alert Monitor	Yes
DHCP status	Yes

■ OpsCenter

Provisioning

- Solaris zones
 - Cryptic names

```
root@opsvml:~# zoneadm list -cv
```

ID	NAME	STATUS	PATH	BRAND	IP
0	global	running	/	solaris	shared
-	opsZ2-migrated	unavailable	/var/mnt/oc-zpools/a205877f-59e4-47f3-8b39ee1c3bc14608/3e52240f-df9a-4035-903b-0c334cb8e110	solaris	excl

- LDoms
 - Not performed due to incompatibility with the latest OVM Manager 3.3 Agent

■ OpsCenter

Migration zones

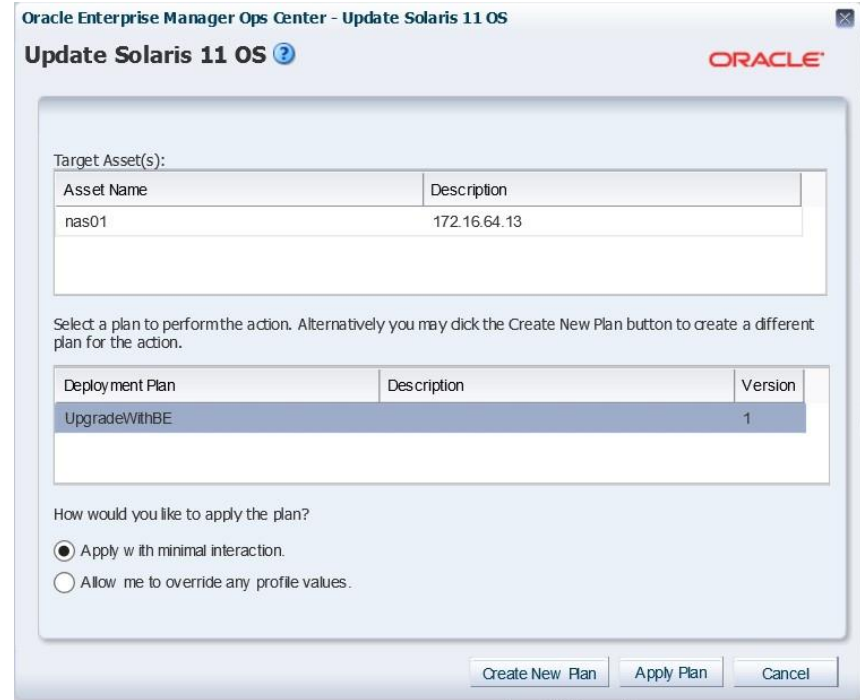
- Server pool with shared storage

```
09/28/2015 03:26:12 PM CEST INFO Waiting for opsZ2 to complete shutdown (72063)
09/28/2015 03:37:42 PM CEST INFO Zone opsZ2 forcibly shutdown (72065)
```

■ OpsCenter

Update

- Nice GUI
- Problems after it



■ OpsCenter

Collaboration with VM manager

- Can be used in combination
- Configuration is going to be applied
- Locking mechanism results

✖ Message

OPS_VM1 is locked. job info: time:1443016443876 name:restartVm-1443016443874
description:OpsCenter: Restarting virtual machine [0004fb00000600001b9541d81708058]
Wed Sep 23 14:58:10 GMT+01:00 2015

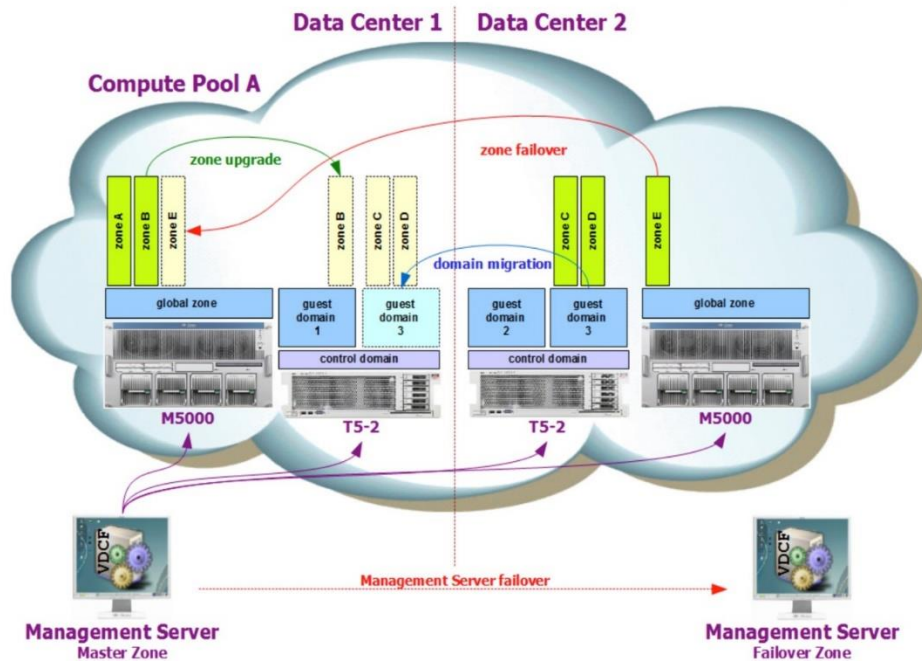
Results

Resultate von OpsCenter	
Support unserer Architektur mit aktuellen Versionen von Solaris und OVM Server	Nein, OVM wird max. in der Version 3.2.X unterstützt
Schmerzfreie Übernahme unserer Bestehenden Konfiguration	Ja
HW Monitoring der SPARC T4 Systeme	ja
Konfigurationsmanagement des SPARC Hypervisors	
Konfigurationsmanagement von Solaris Zonen	Ja, Logins, Netzwerk
Optimale Unterstützung des Netzwerkstacks im Hypervisor und in Solaris	
Provisioning von Solaris und Solaris Zonen	Ja, bei VMs auf x86 ohne Provisioning
Solaris Updates / Patching	Ja
Migration von Zonen / LDoms von einer HW auf die andere	Zonen: Ja, LDOMS:
Monitoring von LDoms, Zonen und Solaris	Ja
Einfaches Setup und Bedienbarkeit der Tools	Ja, UI eher kompliziert

VDCF

VDCF

Overview



■ VDCF

First impression

- commercial product
- Free version
 - Number nodes limited to 5
- No GUI but top command line interface
- Simple architecture
 - Management server
 - Client SW, SSH interfaces
 - Proxy server in need

■ VDCF

Installation

- Simple
 - Install SW components as Solaris packages
- Central configuration file
- User with SSH keys as RBAC configuration
- Command line interface
 - Each command have a good help function

■ VDCF

Takeover of our environment

- «discover» the current environment

```
nodecfg -c discover nonroot hostname=germany.ttc.trivadis.com  
name=germany add
```

- Exclude SAN LUNS's

```
diskadm -c show  
diskadm -c mark name=60002AC00000000000000006B0000D140 foreign  
Deployment
```

■ VDCF

Deployment

- Setup from AI and repositories with ipsadm

```
ipsadm -c show_repo ipsadm  
-c show_build
```

```
ipsadm -c create_service patchlevel=2.14 platform=sparc name=sparc_214  
ipsadm -c create_build name=sparc_214 patchlevel=2.14 service=sparc_214
```

■ VDCF

Deployment

- LDom

```
gdom -c create name=davos cdom=switzerland cores=2 ram=4G comment="Davos
LDOM" diskadm -c mark name=60002AC00000000000000006E0000D140 useable
gdom -c adddisk name=davos type=root guids=60002AC00000000000000006E0000d140
gdom -c addnet name=davos ipaddr=davos.ttc.trivadis.com netmask=255.255.254.0
gdom -c commit name=davos
node -c enable_install name=davos build=sparc_t4_21
```

■ VDCF

Deployment

- Zones

```
vserver -c create name=owl node=bern
vserver -c addnet name=owl type=managment ipaddr=owl.ttc.trivadis.com \
stack=exclusive
dataset -c create name=rpool vserver=owl
\
layout=60002AC00000000000000006C0000D140d0
vserver -c addfs type=root name=owl
vserver -c commit name=owl server -c boot
name=owl
```

■ VDCF

Migration from LDoms and zones

- LDoms

```
cdom -c show gdom
-c show
gdom -c show name=bern candidates
gdom -c migrate name=bern
cdom=germany
```

- Zones

```
node -c show gdom
-c show
vserver -c show name=raven candidates
vserver -c migrate name=raven node=davos
```

■ VDCF

Monitoring

- HW

```
hwmon -c show
```

```

          Current Hardware State
          Node  Model
State  Last Change          Last Update          Console  Soft State  HW
2015-10-21 15:03:39  2015-10-21 15:03:39  ON      ILOM      OS-RUN      OK
          germany  ORCL, SPARC-T4-1
2015-10-21 15:03:39  2015-10-21 15:03:39  ON      ILOM      OS-RUN      OK
          switzerland  ORCL, SPARC-T4-1
```

- Osmon

```
osmon -c show
```

```
No critical osmon data found
```

■ VDCF

Monitoring

- Osmon

```
osmon -c show_fs      [ over=<percent> ]
                        [ summary ]
                        [ root ]
osmon -c show_smf     [ state="state1,state2,state3" ]
                        [ search=<smf name> ]
                        [ server=<server name> ] osmon
-c show_dataset [ over=<percent> ]
                        [ summary ]
                        [ root | dataset ]
osmon -c show_server  server=<server name>
                        [ all ]
osmon -c summary      [ server=<server name> ]
```

■ VDCF

Monitoring

- Ressources

```
rcmon -c show server=bern cpu hourly
rcmon -c show server=davos memory daily
rcmon -c summary node=davos
rcmon -c summary vserver=raven
```

Results

Resultate von VDCF		
Support unserer Architektur mit aktuellen Versionen von Solaris und OVM Server	ja	+++
Schmerzfreie Übernahme unserer Bestehenden Konfiguration	ja	+++
HW Monitoring der SPARC T4 Systeme	ja	+
Konfigurationsmanagement des SPARC Hypervisors	ja	+++
Konfigurationsmanagement von Solaris Zonen	ja	+++
Optimale Unterstützung des Netzwerkstacks im Hypervisor und in Solaris	ja	+++
Provisioning von Solaris und Solaris Zonen	Ja	+++
Solaris Updates / Patching	ja	+++
Migration von Zonen / Ldom's von einer HW auf die andere	ja	+++
Monitoring von LDOM's, Zonen und Solaris	ja	+
Einfaches Setup und Bedienbarkeit der Tools	ja	+++

■ Solaris bordmedium

Building own tools

- Solaris bordmedium
 - «Idm commands»
 - «zonecfg commands»
 - «zoneadm commands»
 - ZFS administration
 - Own scripts
- Puppet ressource
 - Zones

Conclusion

■ Conclusion

Comparison of tools

- VM manager
 - Not yet matured
 - Limitations
- OpsCenter
 - Top for monitoring
 - Complicated setup for virtualization and Solaris components
 - Plan on Oracle support
 - Didn't support the current version from VM manager in summer 2015

■ Conclusion

Comparison of tools

- VDCF
 - Commercial Tool
 - Free version with limited nodes
 - Very easy setup
 - No GUI
 - Top command line interface
 - Meets all our needs

Roman Gächter
Principal Consultant

Roman.gaechter@trivadis .com

