

VDCF - Virtual Datacenter Cloud Framework for the Solaris™ Operating System

Release Notes

Version 8.4
24 October 2023

Copyright © 2005-2023 JomaSoft GmbH
All rights reserved.

1 Introduction

This documentation describes the differences of the releases of the Virtual Datacenter Cloud Framework (VDCF) for the Solaris Operating System. It explains how to upgrade to the newest release.

See these other documents for further information:

<i>VDCF – Installation Solaris 11</i>	for information about installing VDCF on Solaris 11
<i>VDCF – Getting Started</i>	detailed steps if using VDCF for the first time
<i>VDCF – Proxy</i>	for information about running VDCF using Proxies
<i>VDCF – Quick Reference</i>	for a short command overview
<i>VDCF – Administration Guide</i>	for information about the VDCF Usage
<i>VDCF – Resource Management</i>	for information about VDCF Resource Management
<i>VDCF – Monitoring</i>	for information about VDCF Monitoring (HW, Resource, OS)
<i>VDCF – Sun/Solaris Cluster Guide</i>	for information about VDCF Cluster Support

These and all other VDCF documents can be found at:

<https://www.jomasoft.swiss/vdcf/#js-docu>

2 What's new in Release 8.4 (19 October 2023)

2.1 Security/Hardening based on compliance report

Systems can be automatically hardened based on compliance reports. The commands `node -c harden` and `vserver -c harden` offers now the optional 'compliance' flag to fix the failed rules of the last compliance assess.

2.2 OS Monitoring - zpool autoclear

Suspended zpools are automatically cleared by the `osmon` update job. This feature improves application and server availability and is very helpful for short storage outages at night. To enable this feature set `OSMON_ZPOOL_AUTOCLEAR=TRUE`.

2.3 OS Monitoring - Resource Alarming (CPU/RAM)

Resource usage is stored in the VDCF repository using the `rcmon` tool since several years. Now alarming was added to detect and report high resource usage over a defined period.

The default limit is 80% and the default duration period is 15 minutes. The limit can be configured individually for each server where required.

This feature is enabled by default.
Find details about this new feature in the VDCF Monitoring Guide.

2.4 OS Monitoring - Additional recipients per Node

Using `'nodecfg -c modify name=mynode email=team1@customer.com,team2@customer.com'` additional recipients can be defined for each node to receive `osmon` warning update emails.

2.5 VDCF Dashboard (Python 3.11)

The new version of the VDCF Web Dashboard is now based on Python 3.11 and Django 4.1. Existing Dashboards are automatically updated. Solaris 11.4 SRU60 or later includes Python 3.11.

2.6 Pending configuration changes

The new operation `vdcfadm -c show_pending` lists all the pending changes in the VDCF environment. This are all uncommitted dataset, filesystem, network, gdom, vserver and rcadm changes. Using this operation you can detect unfinished configuration work.

2.7 Encrypted zpool dataset

You can now create encrypted zpool datasets with the new 'encrypt' flag of the dataset -c commit operation.

Keystore, Algorithm and Encryption default (ALLDATASET or NONE) is defined with the following configuration variables:

```
DATASET_ZPOOL_ALG          aes-256-ccm
DATASET_ZPOOL_ENCRYPTION  NONE
DATASET_ZPOOL_KEYSTORE    passphrase,file:///root/.zpool.key
```

2.8 Improved import operations

The import operations `node -c import`, `vserver -c import` and `dataset -c import` are improved. You can now also import server and dataset instances which use advanced and unusual configuration and setups options.

VDCF is now able to import LDOMs created by Oracle VM Manager for SPARC and Oracle OpsCenter.

2.9 New SQLite Version 3.41.2 included

This VDCF Release contains and uses the upgraded SQLite Version 3.41.2.

2.10 Performance improvements

Significant performance improvements of the `rcmon -c show` operation

2.11 New operations

<code>osmon -c modify_cpu</code>	To modify threshold of CPU usage alarming for individual servers
<code>osmon -c modify_ram</code>	To modify threshold of RAM usage alarming for individual servers
<code>osmon -c show_cpu</code>	Used to show servers with high CPU usage
<code>osmon -c show_ram</code>	Used to show servers with high RAM usage
<code>vdcfadm -c show_pending</code>	Lists the pending/uncommitted changes in the VDCF env

2.12 Enhanced operations

<code>dataset -c commit</code>	The optional 'encrypt' flag can be used to create zpools with encryption.
<code>ipsadm -c show_build</code>	New 'sparc' and 'i386' flags to limit the builds to the given platform
<code>node -c harden</code>	Optional 'compliance' flag to harden the node based on the last compliance report failed rules.
<code>nodecfg -c modify</code>	New attribute 'email' to specify additional recipients for the VDCF osmon warning/alarming emails.
<code>nodecfg -c modify_net</code>	New flags 'cable/nocable' to mark unused network ports to be ignored by the VDCF operation system monitoring
<code>osmon -c update</code>	Optional 'res' flag to check the resource usage limits (RAM/CPU)
<code>vpool -c show</code>	Additional 'all_server' flag to list all servers including the v pools the servers are defined in.
<code>vserver -c harden</code>	Optional 'compliance' flag to harden the vserver based on the last compliance report failed rules.

2.13 Configuration variables

New settings

CPOOL_ASSIGN_UPDATE_GDOM

If TRUE assigns all Gdoms on a Cdom to the new cpool if a cdom is assigned to a different cpool.

Allowed values: TRUE | FALSE
default: FALSE

GDOM_SHOW_DISK_COMMENT

Shows Disk comments at gdom -c show name= if set to TRUE

Allowed values: TRUE | FALSE
default: FALSE

IPS_REPO_THREADS

Defines the number of threads to use when creating new IPS repo

default: 200

IPS_REPO_UPD_VERIFY

Enables verification of zip files at ipsadm -c update_repo

Allowed values: TRUE | FALSE
default: TRUE

NODE_UPGRADE_FINISH_ENABLE_INSTALL

Defines if a node -c enable_install active should be executed automatically after node -c upgrade_finish to enable the new active build

Allowed values: TRUE | FALSE
default: TRUE

NODE_VERIFY_UPDATE_IPMPGROUP

If set to TRUE node -c verify will update IPMP group names on Solaris 11 to include the interface number in the VDCF repository for higher standardization.

Allowed values: TRUE | FALSE
default: FALSE

OSMON_CPU_DURATION

Duration in minutes while the CPU limit must be reached for osmon to alarm about high usage

default: 15

OSMON_CPU_LIMIT

Default limit in % which must be reached during OSMON_CPU_DURATION minutes for osmon to alarm about high usage.

Individual limits can be defined using `osmon -c modify_cpu server=`

default: 80

OSMON_RAM_DURATION

Duration in minutes while the RAM limit must be reached for osmon to alarm about high usage

default: 15

OSMON_RAM_LIMIT

Default limit in % which must be reached during OSMON_RAM_DURATION minutes for osmon to alarm about high usage.

Individual limits can be defined using `osmon -c modify_ram server=`

default: 80

OSMON_ZPOOL_AUTOCLEAR

`osmon -c update` will automatically executed zpool clear on SUSPENDED zpools if set to TRUE

Allowed values: TRUE | FALSE
default: FALSE

New values of existing settings

NODE_SHOW_ATTR

new values "GROUPPKG" and "LOCATION"

Optional to display the attributes at node -c show

2.14 Deprecated

Deprecated feature “Control Domain on Solaris 10”

It is highly recommended to use Solaris 11 for Control Domains.

VDCF still supports Solaris 10 Control Domains, but the use is deprecated.
Support for Solaris 10 will be removed in a future VDCF Release.

3 How to upgrade VDCF from a previous Release

3.1 VDCF upgrade using bundles

The file name to use depends on your VDCF license (entry, standard or enterprise)

```
$ gunzip vdcf_enterprise_8.4.0_sparc.tar.gz  
$ tar xf vdcf_enterprise_8.4.0_sparc.tar
```

Run the upgrade as root:

```
#!/vdcf_bundle/vdcf_upgrade
```

```
VDCF Package Overview (vdcf / 19.10.2023 08:07:33)
```

Package	Installed	Available	Upgrade
JSvdcf-base	5.5.4	8.4.0	YES
JSvdcf-patch	5.5.4	-	REMOVE
JSvdcf-vserver	5.5.4	-	REMOVE
JSvdcf-ldom	5.5.4	-	REMOVE
JSvdcf-rm	3.2.8	3.6.6	YES
JSvdcf-monitor	2.3.4	3.6.0	YES

```
Missing required Solaris 11 Package shell/expect  
Trying to install ...  
pkg install shell/expect  
Execute VDCF Upgrade (Y/N)? Y
```

```
Removing packages ...  
JSvdcf-monitor      ... done  
JSvdcf-rm           ... done  
JSvdcf-ldom        ... done  
JSvdcf-vserver     ... done  
JSvdcf-patch       ... done  
JSvdcf-base        ... done
```

```
Adding packages ...  
JSvdcf-base        ... done  
JSvdcf-rm          ... done  
JSvdcf-monitor     ... done
```

```
VDCF Upgrade successful.  
Check /var/tmp/vdcf_upgrade.log for details.  
TODO: Upgrade your Nodes using: vdcfadm -c update_node all
```

3.2 Update Client Package

It is **required** to update the VDCF client package on all nodes. If a node is down, please boot it to update the client package.

```
vdcfadm -c update_node all
```

Verify that all nodes have the new VDCF client package installed:

```
vdcfadm -c show_node all
```

4 Known Issues

4.1 VDCF Issues

None

4.2 Open Solaris 11.4 Issues

None

4.3 Fixed Solaris 11 Issues

node -c upgrade

When upgrading from Solaris 11.3 to Solaris 11.4, the default route is lost if it has a route name.

Known Solaris Bug:

29308785 - named default route is not upgraded

Fixed in Oracle Solaris 11.4 SRU 10 / June 2019

Conflict Errors Running pkg update from truncated path: pkg:/laris (Doc ID 2827062.1)

Known Solaris Bug:

Bug 33098570 - non-ascii filenames in pkg causes shortened package fmri and failed upgrade

Bug 31492888 - pkg updates fail with non UTF-8/C locale

looks usually like

These packages cannot be installed together. Any non-conflicting subset of the above packages can be installed.

The following packages all deliver file actions to usr/lib/libraidcfg.so.1:

```
pkg://solaris/system/storage/raid-utilities@11.4,5.11-11.4.44.0.1.113.2:20220324T175551Z
```

```
pkg:/laris/system/storage/raid-utilities@11.4,5.11-11.4.38.0.1.101.4:20210928T224632Z
```

Workaround: remove the files /var/pkg/cache/actions.* in the global zone and associated zones

Fixed in Oracle Solaris 11.4 SRU39 / November 2021

vserver -c apply

```
-bash-4.1$ vserver -c apply name=v0173
```

```
zone 'v0173': warning: FSS scheduling class expected, found instead
```

```
zone 'v0173': Checking: Setting scheduling-class=FSS
```

```
zone 'v0173': Applying the changes
```

Known Solaris Bug:

22935986 - LZR does needless check/apply when scheduling-class is not set

Fixed in Oracle Solaris 11.4 GA / August 2018

5 History: What was new in Release 8.3 (23 August 2022)

5.1 Filesystems of vServer on local directory are now imported and monitored

In older VDCF releases the filesystems of vServers installed on a local directory were ignored. With this release these filesystems are imported and monitored using `osmon`.

5.2 Enhanced Compliance Reporting

Command `osmon -c show_compliance server=` shows the result of each compliance rule.

5.3 Improved and faster Verification Operations

The performance of `node` and `vserver -c verify` is improved.
New command `diskadm -c verify` to report LUNs that are used in multiple cPools.
Command `node -c verify` does now verify aggregated interfaces.

5.4 Support for vServer limitpriv

The default privileges of new vServers can be defined with the config variable `VIRTUAL_LIMITPRIV_DEFAULT`. Existing privileges are displayed by `vserver -c show name=` and can be modified using `vserver -c modify`.

5.5 Additional rules for hardening Node and vServers

VDCF supports now 59 Hardening Rules. See `node -c harden help` for details.

5.6 New 'migration-class2' is supported for SPARC S7/T8/M7/T8/M8 servers

Command `gdom -c show candidates` and `gdom -c migrate` support the new 'migration-class2' for live migration between SPARC S7/T7/M7/T8/M8 Servers. This new migration class includes only CPUs that support the new CPU features like DAX, etc.

5.7 New SQLite Version 3.39.1 included

This VDCF Release contains and uses the upgraded SQLite Version 3.39.1.

5.8 New operations

`diskadm -c verify`

List LUNs which are visible by more than one compute pool

5.9 Enhanced operations

`cpool -c show`

The optional '`sru`' flag will list the systems grouped by Solaris 11 Update and SRU

`dataset -c verify`

Optional '`update_size`' flag will update all detected dataset size changes.

`ipsadm -c update_repo`

Optional '`next`' flag to update an IPS repository to the next available SRU. Optional '`sru`' argument to update a repository with the requested SRU.

`node -c show`
`node -c show name=<node>`

Optional '`cluster`' or '`noncluster`' flag to limit the node list
Optional '`nettype`' argument will list vServer using that nettype.

`osmon -c show_compliance`

Optional '`server`' argument to show compliance details for the requested server

`vdcfadm -c update_node`

Optional '`cluster`' or '`noncluster`' flag to limit the target systems where the JSvdcf-client packages will be replaced.

`vserver -c modify`

Optional '`limitpriv`' argument to modify/update the limitpriv value for a vServer. Optional '`clear_limitpriv`' flag to remove a limitpriv value from a vServer and reset it to the 'default' value.

`vserver -c revert`

Optional '`limitpriv`' flag to revert a limitpriv modification for a vServer to the previous value.

`diskadm -c show`

Optional '`location`' and '`box`' arguments to filter the disk list accordingly.

`diskadm -c modify`

Optional '`location`' and '`box`' arguments to modify disk location and/or Storage Box name for a LUN.

`node/vserver -c verify`

Verify of '`node`' or '`vserver`' will ignore systems that are matching these filter variables:
NODE_VERIFY_IGNORE
VSERVER_VERIFY_IGNORE

5.10 Configuration variables

New settings

VIRTUAL_MACADDR_ENABLE

Enables or disables the import/verify of configured mac addresses for exclusive IP stack

Allowed values: TRUE | FALSE
default: TRUE

VIRTUAL_LIMITPRIV_ENABLE

Enables or disables the import/verify of limitpriv values from the local zone configuration

Allowed values: TRUE | FALSE
default: TRUE

VIRTUAL_LIMITPRIV_DEFAULT

The defined limitpriv settings are deployed to a vServer during installation.

Allowed values: comma separated list of privileges
default: default

DISKS_UPDATE_LOCATION

Specifies if location should be updated for 'all' disks (default), or only for disks with an empty location.

Allowed values: ALL | EMPTY
default: ALL

DISKS_SHOW_BOX

Display additional column 'Box' in the 'diskadm -c show' output.

Allowed values: TRUE | FALSE
default: FALSE

NODE_SHOW_NETTYPE

Define the behavior of 'node -c show name=<node name> nettype=<NETTYPE>'.
With the optional argument 'nettype' all vServer using the requested nettype are listed for that node.

- ALL: list all information about the node, like filesystem, network, patchlevel
- NET: list only general node information and configured network (and related vserver nets)
- NONE: list only general node information (and related vserver nets)

Allowed values: ALL | NET NONE
default: ALL

DATE_FORMAT

Define the format of date output.

Allowed values: The formatting is corresponding to the date command.
default: %d.%m.%Y

REXEC_SSH_CONFIG

Defines the ssh config that should be used by VDCF remote exec user 'vdcfexec'

Allowed values: absolute path to ssh config file
Default: <empty>

NODE_VERIFY_IGNORE

Defines which nodes should not be verified. Goal is to prevent verify errors for nodes with a setup which is not fully supported by VDCF.
A verify for a single node will ignore this variable.

Allowed values: comma separated node list
Default: <empty>

VSERVER_VERIFY_IGNORE

Defines which vServer should not be verified. Goal is to prevent verify errors for vServer with a setup which is not fully supported by VDCF.
A verify for a single vServer will ignore this variable-

Allowed values: comma separated vServer list
Default: <empty>

6 History: What was new in Release 8.2 (23 November 2021)

6.1 Solaris 11 required for VDCF Installation

This new VDCF Release requires Solaris 11. It is not supported to install VDCF on Solaris 10 anymore. Installation, Operation, Management,... of Solaris 10 servers remains supported.

6.2 Monitoring and Alarming

Monitoring of Virtual Memory, new core files and network interface states is now supported by the osmon command.

6.3 Solaris 10 Patching

Because Oracle does not support smpatch (Update Connection) anymore, VDCF patchadm does now use PCA for Solaris 10 Patch Management. It is required to store Oracle User/Password using patchadm -c credentials.

6.4 New SQLite Version 3.35.5 included

This VDCF Release contains and uses the new SQLite Version 3.35.5.

6.5 Data disk remove for zpools is now disabled by default

The operation dataset -c remdisk is disabled for zpools by default. It can be enabled with DATASET_ALLOW_REMDISK=TRUE

We disabled this function because a remove of a disk with lots of data could have major performance impact to read operations. remdisk should only be used to remove an empty disk after it was added to a zpool by mistake.

Details in ZFS Guide Solaris 11.4 / [Removing Devices From a Storage Pool](#)

6.6 node/vserver harden supports additional rules

VDCF supports now 54 Hardening Rules. See node -c harden help for details.

6.7 Performance improvements

Significant performance improvements for osmon and diskadm operations.

6.8 New operations

`osmon -c show_net`

Lists the physical network interfaces including the state (UP/DOWN)

6.9 Enhanced operations

`osmon -c update`

Optional 'net' flag, to update the network interface states only
Optional 'cores' flag, to report new core files only

`osmon -c summary`

New flag 'net' to limit output to critical network interface states

`osmon -c modify_swap`

New flag 'virtualmem' to modify warn limits for swap virtualmemory alarming.

`vserver -c commit`

New flag 'preserve' to preserve the system configuration when installing a branded Zone

6.10 Configuration variables

New settings

CPOOL_PATCH_SYSSHOW

Defines how many systems are listed in the servers column.

Default: 5

DATASET_ALLOW_REMDISK

Enables or disables the dataset -c remdisk operation for zpools.
Using remdisk may have major impact to read performance.

Allowed values: TRUE | FALSE
default: FALSE

DISKS_SHOW_GDOM

Defines which Gdoms are listed at diskadm -c show name=
ALL list all Gdoms
ACTIVE only lists the Gdom where the dataset is active.

Allowed values: ALL or ACTIVE
Default: ALL

NODE_VERIFY_HWBTI

If set node -c verify reports servers with different HWBTI ILOM setting.

Allowed values: empty, "disabled" or "enabled"
Default: empty

OSMON_COREADM_WARN

osmon will report missing directories and coreadm settings.

Allowed values: TRUE | FALSE
default: FALSE

OSMON_CORES_IGNORE

Defines the cores file names which should be ignored.
Useful if applications produce a lot of cores.

Example: export OSMON_CORES_IGNORE=java.400,rlwrap

RESOURCE_NODE_CPU_MIN_THREADS

Defines the number of threads which are reserved for the Global Zone.
To ensure vServers do not get too much CPU threads allocated.

Default: 4

RESOURCE_NODE_RAM_MIN_GB

Defines how much Memory is to be reserved for the Global Zone.
To ensure vServers do not get too much RAM allocated.

Default: 4

6.11 Deprecated

Deprecated feature “Control Domain on Solaris 10”

It is highly recommended to use Solaris 11 for Control Domains.

VDCF still supports Solaris 10 Control Domains, but the use is deprecated.
Support for Solaris 10 will be removed in a future VDCF Release.

7 History: What was new in Release 8.1 (24. March 2020)

7.1 VDCF Dashboard 2.0

The new version of the VDCF Web Dashboard is now based on Python 3.7 and Django 3.1. It includes a new page for datasets. Existing Dashboards are automatically updated. Solaris 11.4 SRU12 or later includes Python 3.7. VDCF Dashboard 2.0 is not supported on older Solaris versions.

7.2 Monitoring

A new feature called "autoclear" in osmon automatically clears SMF services with state 'maintenance'. Free space of filesystems is now monitored and reported.

7.3 ZFS Encryption

New filesystems can be setup with encryption automatically enabled.

7.4 Patching Statistics

cpool -c show provides a new Patching statistic

7.5 Resource Management

rcadm does now support to set the 'Maximum number of Processes' using the PROC attribute for vServer (Solaris zones).

7.6 New SQLite Version 3.32.3 included for Solaris 11

On Solaris 11, this VDCF Release contains and uses the new SQLite Version 3.32.3.

7.7 New operations

`node -c destroy`

Removes a node definition completely.

7.8 Enhanced operations

`cpool -c show`

New flag 'patches [summary]' to view Patching statistic

`gdom -c modify`

Optional 'livemigr' argument to disallow Live Migrations

`node -c console`

New argument 'user' to access the console with a non-default user

`osmon -c modify_fs`

`osmon -c show_fs`

Optional 'warnfree' argument to define an alarm level if a Filesystem reaches the given free limit

`rcadm -c set`

PROC argument to limit the number of processes

`vserver -c attach`

New optional flag 'ignore_upgrade' to ignore the upgrade state of the node

7.9 Configuration variables

New settings

DATASET_ZPOOL_ALG

Defines the algorithm for ZFS encryption

Allowed values: aes-128-ccm,aes-192-ccm,aes-256-ccm,aes-128-gcm,
aes-192-gcm,aes-256-gcm

Default: aes-256-ccm

DATASET_ZPOOL_ENCRYPTION

Defines where to use ZFS encryption

Allowed values: ALLDATASET or NONE

Default: NONE

DATASET_ZPOOL_KEYSTORE

Defines the keystore for ZFS encryption with format:location

Allowed values: format: passphrase,raw,hex

Default: passphrase:file:///root/.zpool.key

OSMON_SMF_AUTOCLEAR

Defines which SMF should be cleared if in state 'maintenance'

Default: svc:/system/filesystem/minimal:default

8 History: What was new in Release 8.0 (24. March 2020)

8.1 Kernel Zones

Kernel Zones are Zones which run an own, independent Solaris Kernel. The Kernel Version can be different from the underlying Solaris 11 Node. Such Kernel Zones can be created using the new `type=KERNEL` argument of the `vServer` command.

RAM and CPUs/cores are dedicated for the Kernel Zones, to make sure they are completely independent of other Zones. Using VDCF the cpu configuration can be changed online. For Memory changes the Kernel Zone must be restarted.

Kernel Zones can be upgraded to another Solaris SRU and migrated live to another Server if the CPUs are the same.

8.2 Dump zpool dataset

Using the new dump flag of the dataset create operation a zpool can be created which is then configured as dump device for nodes or kernel zones.

8.3 Resource Management

The Resource Management feature delivered with the `rcadm` command is now available in the VDCF Free and Entry Edition.

8.4 New SQLite Version 3.28.0 included for Solaris 11

On Solaris 11 this VDCF Release contains and uses the new SQLite Version 3.28.0

8.5 New operations

<code>ipsadm -c modify_build</code>	Modify an existing Solaris 11 Build definition. Currently the AI iso file for Kernel vServer can be modified.
<code>vserver -c upgrade</code> <code>vserver -c upgrade_check</code> <code>vserver -c upgrade_finish</code> <code>vserver -c upgrade_failback</code>	New operations to upgrade Kernel Zones to another Solaris 11 SRU. Including check, finish to cleanup and failback functions.
<code>vserver -c adddisk</code> <code>vserver -c remdisk</code> <code>vserver -c attach_root_mirror</code>	To manage root disks and mirroring of Kernel Zones
<code>vserver -c show_perf</code>	To show a Performance and Resource Usage Overview for a vServer.

8.6 Enhanced operations

<code>dataset -c create</code>	Optional 'dump' argument to create dump zpools for nodes or kernel zones.
<code>diskadm -c show</code>	Optional 'parsable' argument to print machine readable output of a single disk.
<code>vserver -c create</code>	New arguments 'cpus', 'cores', 'ram' and type=KERNEL to create Kernel Zones.
<code>vserver -c modify</code>	New arguments 'cpus', 'cores', and 'ram' to modify the dedicated resources of Kernel Zones.
<code>vserver -c migrate</code>	New optional flag 'live' to start a Kernel vServer live migration.
<code>vserver -c revert</code>	Optional arguments 'resources' and 'rootdisk' to revert pending changes of Kernel Zones.
<code>vserver -c show</code>	Optional argument 'kernel' to limit the vserver list to Kernel Zones.
<code>vserver -c import</code>	Optional 'all' flag to re-import and update all existing vServer.
<code>vpool -c create</code>	New argument 'copyof' to duplicate an existing vpool.

8.7 Configuration variables

New settings

GDOM_CPU_ARGS

Defines the allowed CPU arguments for gdom creation and modification

Allowed values: VCPUS, CORES, MAXCORES
default: VCPUS, CORES, MAXCORES

GDOM_MEMORY_MAX_REDUCE

Defines by how many percent the Memory of a Gdom can be reduced. This variable was introduced to make sure a Gdom is not reduced to much by a typo.

Allowed values: <number> (%)
default: 30

KZONE_MRP

Defines if a Memory Resource Pool (MRP) should be defined for new Kernel Zones.

Allowed values: TRUE | FALSE
default: TRUE

OSMON_DISK_REGISTER

Defines if new Disks detected by osmon should be automatically registered.

Allowed values: TRUE | FALSE
default: TRUE

ROOTDISK_MIN_GB

Defines a minimum root disk size for Gdom and Kernel Zones.

Allowed values: <number> (GB)
default: 10

SERVER_TMP_LIMIT_MB

If this optional VAR is defined, VDCF creates a size limit for the /tmp filesystem for new Servers.

Allowed values: <number> (MB)
default: none

VSERVER_CANDIDATES_LIST_FULL

Defines if vserver -c show candidates should list the full list of candidate nodes. By default (FALSE) only candidates are listed. If TRUE is defined all nodes of the cpool are listed.

Allowed values: TRUE | FALSE
default: FALSE

VSERVER_IMPORT_CHECK_RCADM

By default vserver -c import doesn't allow to import over provisioned vServer. This check can be disabled with setting this VAR to FALSE.

Allowed values: TRUE | FALSE
default: TRUE



New values and defaults of existing Settings

GDOM_SHOW_ATTR new value “LDM”

Optional to display the Ldom Name at `cdom -c show name=`
Recommended if you import existing Ldoms when Ldom Name is different to the hostname.

VIRTUAL_SHOW_ATTR new value “TYPE”

Optional to display the Virtual Type “F=FULL,K=KERNEL,B=BRAND,S=SPARSE”.
Recommended if you use Kernel Zones.

CONFIG_EXEC_COMMAND_DEFAULTUSER new default “nobody”

Earlier VDCF releases used 'root' user as default. This was changed to nobody to make it more
restricted by default.

9 History: What was new in Release 7.2 (25. March 2019)

9.1 Monitoring Enhancements

OS Filesystems

Operating System filesystems (like /, /var, /tmp) are imported into the VDCF Repository at System Install or by using `dataset -c import` for existing Systems. They are imported as Type 'IMMUTABLE' and cannot be changed or removed. The VDCF `osmonitor` reports if the filesystems reach the usage threshold.

Node root dataset and Node Filesystems

Node/GDom root zpools and filesystems are imported into the VDCF Repository at System Install or by using `dataset -c import`. The root zpool is stored as VDCF dataset 'system_root'. Such root datasets and the node filesystems are monitored by the VDCF `osmonitor`.

Monitoring Report

The new operation `'osmon -c show'` produces a Monitoring Report about all open issues, like full datasets and filesystems, SMF with critical states, SAN Disk Path counts below the target definition and hardware issues. Using `osmon -c enable report` a daily Monitoring eMail Report can be activated.

9.2 Node Filesystems Management

Additional data filesystems can now be managed on Node datasets using `node -c addfs, remfs, growfs, revert` and `commit fs`.

9.3 vServer Verification Update

Changes detected by `vserver -c verify` can now be updated in the VDCF repository using the new `update` flag.

9.4 IPS Repo Firmware Updates

SPARC firmware IPS packages can be updated and shown in IPS Repositories by the new `firmware` flag of the `ipsadm -c update_repo` and `show_repo` operations.

9.5 Solaris 11 Node UpgradeState

This Release introduces the Node UpgradeStates 'checked' and 'upgraded'. This state is set by the `node -c upgrade_check` and `upgrade` operations. System changes (for example new vServer, new Dataset,...) which would cause troubles in case of an `upgrade_failback` are not allowed for Nodes with a UpgradeState. The `node -c upgrade_finish` must be used to clear the UpgradeState.

This Feature can be disabled with `NODE_UPGRADE_CHECK_STATUS=FALSE`

9.6 New operations

<code>ipsadm -c check_archive</code>	Check and update Archive infos in the VDCF Repository
<code>node -c addfs</code> <code>node -c remfs</code> <code>node -c growfs</code> <code>node -c revert</code> <code>node -c commit</code>	New operations for Node Filesystems Management
<code>osmon -c show</code>	Monitoring Report about all critical objects Dataset, Filesystems, SMF, SWAP, DiskPath and Hardware

9.7 Enhanced operations

<code>dataset -c import</code>	Additionally imports the OS and node filesystems. New flag 'all' to import the datasets and filesystems from all nodes.
<code>dataset -c show</code> dataset in combination with the <code>name</code> argument.	Flags 'parsable' and 'header' can now be used for single dataset.
<code>dataset -c verify</code>	Use the 'no_cdom_discover' flag to skip discovery of the underlying CDMs. This makes sense if your CDMs were recently discovered before. Skipping CDM discovery makes the verify process a lot faster.
<code>diskadm -c show</code>	Optional 'size' argument to limit the list of disks with the given size.
<code>diskadm -c update</code>	Optional <code>name</code> argument to update only the specified disks.
<code>gdom -c show</code>	Optional flags 'allfs' and 'datafs' to limit the listed filesystems
<code>hwmon -c show</code>	Optional 'full' flag to list the complete state history of a node. Default is to list <code>HWMON_SHOW_LINES</code> number of lines.
<code>ipsadm -c show_repo</code>	Optional 'firmware' flag to list only the firmware IPS packages.
<code>ipsadm -c update_repo</code>	Optional 'firmware' flag to update firmware IPS packages only.
<code>node -c enable_install</code>	New function 'all active' does re-enable all Nodes which are not currently enabled with the effective active build.
<code>node -c import</code>	Optional arguments 'hostname' and 'proxy'

<code>node -c show</code>	Optional flags 'allfs' and 'datafs' to limit the listed filesystems of a Node. And new flags 'upgraded' and 'checked' to list nodes in the given Solaris 11 Upgrade states.
<code>node -c upgrade</code>	Upgrade now sets the Node to 'upgraded' state to avoid changes on the system.
<code>node -c upgrade_check</code>	Upgrade check now sets the Node to 'checked' state to avoid changes on the system. New 'clear' flag to remove the 'checked' state.
<code>node -c verify</code>	Additional components and configurations are verified: Node filesystems, network interface MAC addresses, network interface linknames, Boot disks New 'nodiscover' flag to avoid node discovery, useful if verify is re-run after changes in the VDCF repository.
<code>nodecfg -c modify linknames from the nodecfg.</code>	New flag 'clear_linknames' to remove network interface
<code>osmon -c disable</code>	'report' flag to disable the OS Monitoring report cronjob
<code>osmon -c enable</code>	'report' flag to enable the OS Monitoring report cronjob
<code>osmon -c modify_fs</code>	New 'server' argument replaces the previous 'vserver' argument because Node filesystems are also monitored.
<code>osmon -c show_fs</code>	New 'parsable' and 'header' flags for machine parsable output.
<code>osmon -c update</code>	Reports details for failed SMF services
<code>rcadm -c show</code>	New 'cpool' argument to list servers of the compute pool
<code>serverconfig -c exec</code>	New 'servertype' argument to execute on all systems of the given type: all, node, vserver, gdom or cdom.
<code>vdcfadm -c statistics</code>	Reports how many Solaris 10 systems are still existing. Migration to Solaris 11 is highly recommended.
<code>vserver -c import</code>	'name' argument to re-import vServer which are already in the VDCF repository.
<code>vserver -c renamefs</code>	'revertonerror' flag to revert the filesystem rename in the VDCF repository in the case the rename on the Node fails.
<code>vserver -c verify</code>	'update' flag to update the VDCF repository with the differences found (filesystems and resources)

9.8 Configuration variables

New settings

HWMON_SHOW_LINES

Defines how many lines are listed by `hwmon -c show node=`

Allowed values: <number>
default: 20

NODE_FILESYS_SHOW

Defines which filesystem are show by default

Allowed values: all | data | none
default: all

NODE_UPGRADE_CHECK_STATUS

Defines if the UpgradeStates of Solaris 11 Nodes are checked when adding, attaching or detaching vservers and datasets.

Allowed values: TRUE | FALSE
default: TRUE

OSMON_REPORT_INTERVAL

Defines at what time the osmon report cronjob should be executed if enabled by `osmon -c enable report`

Default: 0 8 * * 1-5 (weekdays mornings 08:00)

10 History: What was new in Release 7.1

(4. September 2018)

10.1 Oracle Solaris 11.4 supported

This VDCF Release supports the recently released Solaris 11 Update 4.

Delegated zpools are added and removed to vServers without downtime using dataset -c commit
Route names can be specified at routecfg
Disk can be removed from zpools with dataset -c remdisk

10.2 Oracle VM for SPARC / LDom 3.6 supported

This VDCF Release supports the current release of LDom Version 3.6.

10.3 Oracle SPARC T8/M8 Servers Support

This VDCF Release supports the Oracle SPARC T8/M8 Servers

10.4 ZPOOL Online Replication

A zpool can be replicated online to another zpool using the new dataset operations replicate and activate_replica. Using this replication a zpool can be replaced with a smaller zpool. After replication the zpool fragmentation is reduced.

10.5 vServer Verification

Configurations of running vServers (Solaris Zones) can be compared to the data in the VDCF Repository using the 'verify' operation. The verify operation compares filesystem, network and resource settings.

This operation helps to detect and update the VDCF Repository after manual changes were made on the systems..

10.6 Visible VLAN Verification

The VLANs visible on the Node Network interfaces can be imported using node -c update vlan. When adding network configurations and when migrating systems the available VLANs are checked. cpool -c check reports if Nodes in the same cpool don't have the same VLANs visible.

10.7 DNS Client Configuration Updates

Using `type=DNS` at `serverconfig -c exec` you can update your DNS client configuration on your systems very efficient.

10.8 Customer and System specific IPS packages and SMF configurations

New AI configuration directories and files allow automated installation of IPS packages and SMF manifest files. See chapter 2.14 for details.

10.9 Patch Level requirements checked for SPARC S7,M7 and M8 CPUs

The new SPARC CPUs require a minimum Solaris Patch Level. This Patch Level is check by `gdom -c attach`, `node -c enable_install` and `flash -c enable_install`.

10.10 Conservative ZPOOL Mirroring

By default `dataset -c attach_mirror/commit` does attach all disks at once. For very large zpools with many disks this may produce too much load on the target system.

You can define how many disk should be attached at once (`DATASET_ATTACH_MIRROR_MAX_DISK`). Additional disks are then attached in the background.

Define at dataset/zpool size, when this feature should be activated. (`DATASET_ATTACH_MIRROR_MAX_DISK_MINSIZE_GB`)

10.11 New operations

<code>dataset -c replicate</code>	Initiate a zpool replication
<code>dataset -c activate_replica</code>	Activate the replicated zpool
<code>dataset -c cancel_replication</code>	Cancel a replication and cleanup temporary snapshots
<code>vserver -c verify</code>	Verify zone filesystem, network and resource configuration against the VDCF repository

10.12 Enhanced operations

<code>cdom -c show</code>	new flags 'u1, u2, u3, u4' to list only the relevant S11 update
<code>console -c modify</code>	new argument 'proxy' to define a different proxy to access the console
<code>dataset -c commit</code>	new flag 'upgrade' to upgrade the zpool version when removing disks.
<code>dataset -c remdisk</code>	on Solaris 11.4 it is supported to remove disks from a zpool.
<code>gdom -c show</code>	new flags 'u1, u2, u3, u4' to list only the relevant S11 update new flag 'relevant' to list only real candidates, if GDOM_CANDIDATES_LIST_FULL is set to TRUE.
<code>gdom -c modify</code>	new argument 'locked' to restrict gdom migration completely (true) or only in the datacenter location (location).
<code>ipsadm -c show_build</code>	new flags 'u1, u2, u3, u4' to list only the relevant S11 update
<code>node -c reboot</code>	new flag 'wait' to wait till the reboot is finished (helpful for scripts)
<code>node -c show</code>	new flags 'u1, u2, u3, u4' to list only the relevant S11 update new flag 'physical' to list only physical nodes (hides gdoms)
<code>node -c update</code>	new argument 'cpool' to update all nodes in the given cpool new flag 'vlan' to import the vlans visible on the network interfaces.
<code>node -c upgrade</code>	new arguments 'cluster_version' and 'geo_version' to specify the version when upgrading Cluster Nodes. new flag 'wait' to wait till the reboot is finished (helpful for scripts)
<code>node -c upgrade_check</code>	new arguments 'cluster_version' and 'geo_version' to specify the version when upgrading Cluster Nodes.
<code>node -c upgrade_failback</code>	new flag 'wait' to wait till the reboot is finished (helpful for scripts)
<code>node -c upgrade_finish</code>	new flag 'keep' to keep the previous BootEnvironment.

<code>nodecfg -c modify</code>	new attribute 'scheduler' to define the scheduler class before installing a node. The global default is defined as <code>NODE_SCHEDULER_DEFAULT</code>
<code>nodecfg -c modify_net</code>	new flags 'standby/clear_standby' to define the IPMP interface config
<code>nodecfg -c show</code>	new flag 'physical' to list only physical nodes (hides gdoms)
<code>routecfg -c add</code>	new optional argument 'name' for new routes on Solaris 11.4
<code>routecfg -c remove</code>	new argument 'name' to identify the route to be removed.
<code>routecfg -c verify</code>	new argument 'vserver' to verify a single vserver.
<code>serverconfig -c exec</code>	new flags 'parsable/quiet' to reduce output for scripts when executing commands. type=DNS is supported to distribute DNS Client configs.
<code>vserver -c addnet</code>	new optional argument 'probes' to specify probe addresses for IPMP interfaces
<code>vserver -c show</code>	new flags 'u1, u2, u3, u4' to list only the relevant S11 update

10.13 Configuration variables

New settings

DATASET_ATTACH_MIRROR_MAX_DISK

Defines how many disks are attached to a zpool when using `dataset -c attach_mirror/commit`.
By default all disks are attached.

If `DATASET_ATTACH_MIRROR_MAX_DISK_MINSIZE_GB` is defined, the feature is only activated for zpools larger than the `MINSIZE_GB`.

Allowed values: <number>
default: <empty>

DATASET_ATTACH_MIRROR_MAX_DISK_MINSIZE_GB

Defines the size of a zpool when the conservative zpool mirroring feature should be activated.
This variable is only taken into account if `DATASET_ATTACH_MIRROR_MAX_DISK` is also defined.

Allowed values: <number>
default: <empty>

GDOM_VIRTUAL_NET_VERBOSE

Defines if all the details of a network interface should be displayed.
If set to `TRUE` the output looks like `"net0->vnet0->i40e0"`

Allowed values: `TRUE | FALSE`
default: `FALSE`

GDOM_CANDIDATES_ATTR

Define attributes to be shown in `gdom -c show candidates`

Allowed values: <empty> | "LOCATION"
default: <empty>

IPS_REPO_UPD_ASK

Defines if `ipsadm -c update_repo` should ask for confirmation.
`NO` is helpful if `update_repo` is used in a script.

Allowed values: `YES | NO`
default: `YES`

NODE_SHOW_ATTR

Define attributes to be shown in node -c show list.

Allowed values: <empty> | "BE"
default: <empty>

NODE_SCHEDULER_DEFAULT

Defines the default scheduling class for new nodes.

Allowed values: <empty> | "FSS" | "TS"
default: <empty>

SC_LDOM_LOAD_FACTORS

Only relevant for Solaris Cluster Guest Domains. Defines if Load Factors for Memory and CPU should be configured for the Cluster GDoms. Recommended for Clusters with more than 2 Cluster Nodes.

Allowed values: TRUE | FALSE
default: FALSE

VIRTUAL_FS_RETRIES

Defines the number of retries for ZFS filesystem operations.
Taken into account by vserver -c remfs/commit, mount and unmount.

Helpful for very busy zpools.

Allowed values: <number>
default: 3

10.14 New configuration files

ips-pkg.list

Defines IPS packages to install on new Solaris 11 systems.

For IPS packages to install on all nodes:
`/var/opt/jomasoft/vdcf/ai/all-nodes/ips-pkg.list`

For IPS packages to install on all vservers:
`/var/opt/jomasoft/vdcf/ai/all-vservers/ips-pkg.list`

For IPS packages to install on individual servers
`/var/opt/jomasoft/vdcf/ai/<server>/ips-pkg.list`

Sample:
`pkg://solaris/system/management/puppet`
`pkg://solaris/service/network/ntp`
`pkg://solaris/diagnostic/top`

System configuration profiles

All xml files in the following directories will enable additional smf services at installation time.

`/var/opt/jomasoft/vdcf/ai/all-nodes/*.xml`
`/var/opt/jomasoft/vdcf/ai/all-vservers/*.xml`
`/var/opt/jomasoft/vdcf/ai/<server>/*.xml`

Sample to enable puppet:

```
# cat /var/opt/jomasoft/vdcf/ai/node99/puppet.xml

<service name='application/puppet' version='1' type='service'>
  <instance name='master' enabled='true'>
    <property_group name='start' type='application'>
      <propval name='timeout_seconds' type='count' value='10' />
    </property_group>
  </instance>
</service>
```

11 History: What was new in Release 7.0 (25. October 2017)

11.1 Fujitsu SPARC M12 Servers Support

This VDCF Release supports the Fujitsu SPARC M12 Servers.

11.2 Oracle VM for SPARC / LDom 3.5 supported

This VDCF Release supports the current release of LDom Version 3.5.

11.3 VDCF Web Dashboard

With this Release a Browser based Dashboard is available to list general information about VDCF and its objects, including the new Security Compliance Reports.

This Feature is available for Standard and Enterprise Customers when running VDCF on Solaris 11.3.

11.4 Solaris 11.3 Security Compliance Assessment and Reporting

Solaris Security Compliance Reporting is available for Solaris 11.3 Nodes and vServer.

11.5 Solaris 11 Hardening

Hardening for Solaris 11 Nodes and vServer can be used to secure your infrastructure based on Hardening Profiles.

11.6 Network Route Configuration / Verification

On Solaris 11 static routes can be imported and verified against the VDCF Repository. The routes are configured during a Re-Installation. Routes can be added and removed using the new `routecfg` command.

11.7 MPxIO Disk Path Monitoring

With this Release the OS Monitor (`osmon`) has been extended with the MPxIO Disk Path monitoring. The 'Current Path Count' will be checked against the 'Target Path Count'.

11.8 ZPOOL Dataset Log Disk Support

For ZPOOL datasets Log disks can be added and removed. Existing ZPOOL datasets with Log disks can be imported into the VDCF Repository using the `'dataset -c verify'` command.

11.9 Solaris 11 'node -c upgrade_finish' zpool upgrade

When upgrading Solaris 11 Nodes using `node -c upgrade` a new BootEnvironment is created. Using the '`node -c upgrade_finish`' the old BE will be removed and the ZPOOL of the node can be upgraded to the latest ZPOOL version. This features is disabled by default. (NODE_UPGRADE_FINISH_ZPOOL)

11.10 Node and vServer rpool compression

For Solaris 11.3 systems the compression for the rpool can be enabled to save storage space. (NODE_ROOT_POOL_COMPRESSION and VSERVER_ROOT_DATASET_COMPRESSION)

11.11 Clear .ssh/known_hosts on VDCF

If a system is reinstalled the SSH hostkey changes. To prevent the key-mismatch error message, VDCF can remove the old hostkey entry from the known_hosts files. (INSTALL_CLEAR_KNOWN_HOSTS)

11.12 New commands

`routecfg` New command to configure/manage persistent routes on Solaris 11.
Import, verify, deploy, add and remove routes for existing server.

`routecfg -c import` `node=<node name> | all`

`routecfg -c verify` `node=<node name> | all`

`routecfg -c show` [`node=<node name> [full] | vserver=<vserver name>]`
[`destination=<address[/prefix]>]`
[`gateway=<address>]`
[`parsable [header]]`

`routecfg -c add` `destination=<address[/prefix]> gateway=<address>`
`node=<node name> | vserver=<vserver name>`

`routecfg -c remove` `destination=<address[/prefix]> | destination=all`
[`gateway=<address>]`
`node=<node name> | vserver=<vserver name>`

`routecfg -c revert` `node=<node name> | vserver=<vserver name>`

`routecfg -c commit` `node=<node name> | vserver=<vserver name>`

`routecfg -c diff` `server=<vserver1,node2>`

11.13 New operations

`osmon -c assess` Run Compliance Assessment on a node or vserver
`node -c assess`
`vserver -c assess`

`node -c harden` Run a hardening profile on a node or vserver
`vserver -c harden`

`osmon -c show_disk` List the State of Disks Path Monitoring

`osmon -c modify_disk` Modify the target path count for one or more disks of a node to reflect the actual SAN configuration.

`dataset -c addlog` Add one or more Log Disk to an existing ZPOOL dataset.

`dataset -c remlog` Remove one or more Log Disk from an existing ZPOOL dataset.

`dataset -c remdisk` Remove one or more RAW Devices from a RAW dataset.

11.14 Enhanced operations

<code>node -c boot</code> <code>node -c reboot</code> <code>node -c shutdown</code>	now a node list is allowed
<code>node -c import</code>	new flag 'nodeonly' to not import vServer from node.
<code>node -c upgrade_finish</code>	upgrades all ZPOOLS if NODE_UPGRADE_FINISH_ZPOOL is set to TRUE.
<code>nodecfg -c modify</code>	new argument 'benchmark' to define a specific security assessment benchmark for the Node
<code>gdom -c migrate</code>	new flag 'norescheck' to ignore the resource check for guest domains. Only supported on Cluster Control Domains.
<code>osmon -c update</code>	new flag 'disk' will check the MPxIO disk path count
<code>osmon -c summary</code>	new flag 'disk' to show MPxIO disk path count summary
<code>vserver -c renamefs</code>	new flag 'keepfs' to flag to only rename the mountpoint and not the zfs filesystem name.
<code>vserver -c clonefs</code>	new argument 'filesystem' to directly clone a single filesystem or snapshot
<code>vserver -c modify</code>	new argument 'benchmark' to define a specific security assessment benchmark for the vserver.

11.15 Configuration variables

New settings

DISK_DEFAULT_PATH_COUNT

Defines the default MPxIO Disk path count used to compare new SAN Disk against.

Allowed values: <number>
default: <empty>

INSTALL_CLEAR_KNOWN_HOSTS

Defines if the .ssh/known_hosts files should be updated if a system is reinstalled.

Allowed values: TRUE | FALSE
default: FALSE

INSTALL_USE_ROUTECFG

Defines if routes should be used from routecfg instead for config ROUTE/DEFAULTROUTE when installing Nodes and vServers.

Allowed values: TRUE | FALSE
default: FALSE

NODE_ROOT_POOL_COMPRESSION

When set compression is configured on the node rpool.

Allowed values: <empty> | on
default: <empty>

NODE_UPGRADE_FINISH_ZPOOL

Upgrades all ZPOOLS at node -c upgrade_finish if set to TRUE.

Allowed values: TRUE | FALSE
default: FALSE

VSERVER_ROOT_DATASET_COMPRESSION

When set, compression is configured on vServer root dataset. This setting is used on Solaris 11.3 only.

Allowed values: <empty> | lz4
default: <empty>

11.16 Deprecated

Deprecated feature “Control Domain on Solaris 10”

It is highly recommended to use Solaris 11 for Control Domains.

VDCF still supports Solaris 10 Control Domains but the use is deprecated.
Support for Solaris 10 will be removed in a future VDCF Release.

12 History: What was new in Release 6.0 (25. April 2017)

12.1 New SQLite Version 3.11.1 included for Solaris 11

On Solaris 11 this VDCF Release contains and uses the new SQLite Version 3.11.1. Compared to the replaced SQLite Version 3.7.2 we see major performance improvements. The old prerequisite SMCsqlite Package is not required and used anymore.

12.2 Import of existing LDoms

Import an existing Guest Domain into the VDCF database. The import discovers the Guest Domain, Packages, Patches, Disks, Dataset and vServers. Imported Guest Domains can be used the same like the ones created with VDCF.

12.3 Shutdown Feature added to High Availability / Node Evacuation

The VDCF High Availability Solution (hamon) evacuates vServers from a Node based on the Resource Requirements. To make this Evacuation work Nodes with Free Resources are required.

In a two Node Environment with one Test and one Production Server, there are typically not enough Resources available on the Test Environment if the Production Server fails. The new shutdown Feature solves this by shutting down vServers on the Target Nodes, based on pre-configured Categories and Priorities to free Resources.

12.4 SWAP Usage Monitoring

With this Release the OS Monitor (osmon) has been extended with the SWAP usage monitoring. The usage of the SWAP devices based on 'swap -l' listing are monitored.

12.5 Power Usage Statistic

With this Release the HW Monitor (hwmon) collects the power usage for each physical server. The summary can be listed with the new operation 'hwmon -c show_power'

12.6 Readonly Guest Domains

New READONLY State for Guest Domains, which blocks certain commands to prevent configuration changes to the domain.

12.7 System Verification Operations

Configurations on running Systems can be compared to the data in the VDCF Repository using 'verify' operations. The verify operations are available for the node and dataset command. This two operations help to detect and update the VDCF Repository after manual changes.

12.8 Dataset new ZFS Volumes

ZPOOL datasets can be deployed on dynamically created ZFS Volumes (ZVOL) using the newzvol flag. This makes Deployment easier on Systems without access to Shared Storage.

12.9 vServer ZFS Cloning and shared ZPOOLS

vServer ZFS data filesystems can now be cloned and mounted into another vServer on the same Node using the new clonefs operation. Very useful for example for large databases to create a "copy" for Testing or Verification purposes. After cloning the ZPOOL dataset is shared between the vServers.

12.10 Solaris 11 Upgrade Customer Script

When upgrading Solaris 11 Nodes using `node -c upgrade` a new BootEnvironment is created. Applications which scan or modify all filesystems may make troubles. This Release adds a Customer script which can be executed before and after an upgrade to shutdown/start such Applications.

You need to place your individual script at `/var/opt/jomasoft/vdcf/conf/node_upgrade.customer`

12.11 Configure SNMP ASR on ILOM, ILOMx86

VDCF can automatically configure SNMP on your ILOM or ILOMx86 console, to integrate the ILOM into your ASR (Oracle®, Auto Service Request) infrastructure.

The SNMP configuration is done when you execute `console -c add`.

You need to configure your ASR Manager using the VDCF `ASR_MANAGER` variable.

12.12 New operations

<code>dataset -c verify</code>	Verify ZPOOL dataset configuration against the Solaris global zone
<code>hwmon -c show_power</code>	List the power usage for each physical server
<code>ipsadm -c rebuild_repo</code>	Rebuilds the search index of an existing repository
<code>node -c verify</code>	Verify node network configuration against the Solaris global zone
<code>osmon -c modify_swap</code>	Allows to set individual warnover limits for swap monitoring
<code>osmon -c show_swap</code>	Displays the state of swap Monitoring
<code>vserver -c clonefs</code>	Used to clone and mount ZFS filesystems for vServer running on the same Node.
<code>vserver -c make_exclusive</code>	Converts vServer shared ip-stack network to exclusive ip-stack
<code>zfsadm -c rename</code>	Renames existing vServer ZFS snapshots

12.13 Enhanced operations

<code>config -c show</code>	New flags 'os' and 'platform' for selective show
<code>dataset -c create</code>	New flags 'newzvol' and 'zpool' argument to dynamically create new ZFS Volumes.
<code>gdom -c modify</code>	New argument 'readonly' to enable/disable READONLY State for Guest Domains.
<code>gdom -c show</code>	New flags 'active' and 'all-states'. 'active' will only show ACTIVE (cState) guest domains
<code>ipsadm -c create_repo</code>	New optional argument 'zpool' to specify where to store the new repository.
<code>ipsadm -c update_repo</code>	New argument 'p5pfile' to update repository with package file or IDR
<code>node -c enable_install</code>	Use the new 'active' flag to enable the current ActiveBuild.
<code>node -c evacuate</code>	Use the new 'shutdown' flag to shutdown vServer on the target Nodes to free resources, if required.
<code>node -c import</code>	Now supports to import existing Guest Domains when installed on MPXIO disks.
<code>node -c install</code>	Use the new 'wait' flag to wait till the Node Installation is completed. After that you can deploy vServers.
<code>osmon -c summary</code>	New flag 'swap' to display the summary about swap Monitoring
<code>vserver -c show</code>	New flags 'active' and 'all-states'. 'active' will only show ACTIVATED (cState) vServer

12.14 Configuration variables

New settings

ASR_MANAGER

Defines where the ASR manager is running

Allowed values: VDCF_MNGT | VDCF_PROXY | <IP-Address>
default: <empty>

ASR_ILOM_VALUES

Defines the ASR SNMP configuration on the ILOM system controller

default: 1:162:2c:public

rule to configure on ILOM (/SP/alertmngt/rules/) 1 – 15
destination port on ASR Manager
snmp version 1, 2c and 3 allowed
community name or username if version 3 is used

ASR_ENABLE_TEST_RULE

Defines if the service processor should create a test service request after ASR on ILOM configuration.

Allowed values: TRUE | FALSE
default: TRUE

DATASET_ZPOOL_COMPRESSION

When set, compression is configured on new ZPOOLS. This setting is used on Solaris 11.3 only.

Allowed values: <empty> | lz4
default: <empty>

DEFAULT_SIZE

Defines how sizes of Memory, Disks, Datasets and Filesystems are displayed.

Allowed values: GB | MB
default: GB

DISKS_NEWZVOL_POOL

Defines the default ZPOOL used to create ZFS Volumes when using the 'newzvol' flag at 'dataset -c create'. Can be overwritten when using the 'zpool' argument.

default: rpool

GDOM_CANDIDATES_LIST_FULL

Defines if 'gdom -c show candidates' should list all gdom from cpool.

Allowed values: TRUE | FALSE
default: FALSE

GDOM_COMMENT_REQUIRED

Defines if 'gdom -c create' requires the 'comment' argument.

Allowed values: TRUE | FALSE
default: FALSE

GDOM_IMPORT_SET_READONLY

Defines if 'node -c import' should set the guest domains state to READONLY after import.

Allowed values: TRUE | FALSE
default: FALSE

GDOM_SHOW_STATES

Defines if guest domains with active (ACTIVE) state should be listed or all states.

Allowed values: ACTIVE | ALL-STATES
default: ALL-STATES

NODE_UPGRADE_CHECK_GDOM

Defines if 'node -c upgrade' should check if guest domains are on the node (control domain).

Allowed values: TRUE | FALSE
default: TRUE

NODE_UPGRADE_CHECK_VSERVER

Defines if 'node -c upgrade' should check if vServer are on the node.

Allowed values: TRUE | FALSE
default: TRUE

OSMON_SWAP_WARNING

If this SWAP usage (in %) is reached osmon will send a WARNING eMail.

default: 60

REXEC_SSH_QUIET

Defines if ssh remote requests should use the quiet flag (-q).

Allowed values: TRUE | FALSE
default: FALSE

VSERVER_COMMENT_REQUIRED

Defines if 'vserver -c create' requires the 'comment' argument.

Allowed values: TRUE | FALSE
default: FALSE

VIRTUAL_EVACUATION_SHUTDOWN_CATEGORIES

Defines vServer of which Categories should be shutdown by hamon if not enough free Resources are available on Target Nodes.

node -c evacuate takes this Definition into account when the 'shutdown' flag is used.

Multiple Categories must be separated by comma.

VIRTUAL_SHOW_STATES

Defines if vServer with active (ACTIVATED) state should be listed or all states.

Allowed values: ACTIVE | ALL-STATES
default: ALL-STATES

New default values

VIRTUAL_NETSTACK_DEFAULT_S11

Sets the default network stack for vServer on Solaris 11

Allowed values: SHARED | EXCLUSIVE | PRIVATE
default: **EXCLUSIVE** (was SHARED in previous Releases)

12.15 Deprecated

Deprecated feature “Control Domain on Solaris 10”

It is highly recommended to use Solaris 11 for Control Domains.

VDCF still supports Solaris 10 Control Domains but the use is deprecated.
Support for Solaris 10 will be removed in a future VDCF Release.

12.16 Removed Support, Operations and Flags

Removed Support for Platform x86 32-bit

This VDCF Releases does not support 32-bit x86 Nodes anymore.

Removed Operations

`osmon -c show`

Replaced by `osmon -c show_fs`

Removed Flags

`node -c inactivate force`

`gdom -c addnet ipmp`

Addnet automatically uses IPMP if the Control Domain or IO Domain Network has IPMP configured.

13 History: What was new in Release 5.7 (31. March 2016)

13.1 Fujitsu SPARC M10 Servers Support

This VDCF Release supports the Fujitsu SPARC M10 Servers

13.2 Enhanced Solaris 11 Upgrades

Multiple Nodes can be upgraded in parallel. Additional operations to check, failback and clean Boot Environments.

13.3 Extended OS Monitoring

With this Release the OS Monitor (osmon) allows to set individual thresholds for datasets and filesystems.

13.4 Secure Zones with Support for Immutable Zones

It's now possible to change the file-mac-profile zonecfg property using the vserver -c modify command. With this you can easily set the root filesystem of a vServer to be immutable (read-only).

13.5 VDCF Delivery Changes

The previous 3 packages JSvdcf-patch, JSvdcf-vserver and JSvdcf-ldom are integrated into JSvdcf-base. The upgrade process will remove the 3 deprecated packages automatically.

Previous VDCF vServer and LDom Administration Guides are now integrated into the VDCF Administration Guide. You find all information in one place now.

13.6 New operations

<code>node -c import</code>	Import an existing physical node into the VDCF database. This import discovers the Node, Control Domain, Packages, Patches, Disks, Datasets and vServers.
<code>node -c upgrade_check</code>	Trial-Run of Upgrade to the specified Solaris 11 build.
<code>node -c upgrade_failback</code>	Failback to previous Boot Environment.
<code>node -c upgrade_finish</code>	Clears the previous Boot Environment.
<code>osmon -c modify_fs</code>	To set or clear an individual threshold for a filesystem.
<code>osmon -c modify_dataset</code>	To set or clear an individual threshold for a dataset.
<code>vserver -c apply</code>	Applies/Activates the current zonecfg for a running vServer on Solaris 11.2 or later.

13.7 Enhanced operations

<code>diskadm -c modify</code>	New argument 'location' to modify the disk location.
<code>osmon -c show_fs</code>	New flag 'warnover' to list filesystems with an individual threshold set.
<code>osmon -c show_dataset</code>	New flag 'warnover' to list datasets with an individual threshold set.
<code>vserver -c modify</code>	Use the 'file-mac-profile' argument to set the corresponding zonecfg property for installed Solaris 11 vServer. This property is used to control the immutable zones feature of Solaris 11 (secure and read-only). Read the VDCF Administration Guide for more details.

13.8 Configuration variables

New settings

AI_DEFAULT_SERVICE_PLATFORM

Defines the default platform when creating Solaris 11 Automated Installer Services. Customers installing i386 Nodes should set this variable to 'i386'.

Allowed values:	sparc i386
default:	sparc

HWMON_EVENT

Enable/Disable the monitor event alarming for hardware monitoring. This variable replaces the former MONITOR_EVENT variable.

allowed:	true/false
default:	true

IPS_REPO_UPD_AI

Defines if 'ipsadm -c update' should load AI Service packages into the IPS Repository.

Allowed values:	TRUE FALSE
default:	TRUE

IPS_REPO_UPD_FIRMWARE

Defines if 'ipsadm -c update' should load Firmware packages into the IPS Repository.

Allowed values:	TRUE FALSE
default:	FALSE

NODE_UPGRADE_FREE_GB

Defines the required minimum free disk space on the Node and vServer root pools to allow 'node -c upgrade' to run.

default: 10

OSMON_ZPOOL_STATE_OF_INTEREST

Defines the ZPOOL states for which a WARNING eMail should be sent.
By default the critical states are defined.

The 'RESILVERING' state can be added to additionally get eMails when a mirror operation starts and ends.

allowed: FAULTED|DEGRADED|SUSPENDED|RESILVERING
default: FAULTED|DEGRADED|SUSPENDED

OSMON_EVENT

Enable/Disable the monitor event alarming for OS monitoring. This variable replaces the former MONITOR_EVENT variable.

allowed: true/false
default: true

SOL11_UPDATE_FOR_REPO_SYNC

Defines which Solaris11 Update to load into the IPS Repository. This variable is used to save time and avoid to load old Updates into the Repository. Used by 'ipsadm -c update'

Example: `export SOL11_UPDATE_REPO_SYNC=2`

default: 3

VIRTUAL_SHOW_ATTR

Define attributes to be shown in vserver show detail view:

Example: `export VIRTUAL_SHOW_ATTR=WARNOVER`

allowed values: WARNOVER or empty string
default: WARNOVER

13.9 Deprecated

Deprecated platform x86 32-bit

Future VDCF Releases will not support 32-bit x86 Nodes anymore.

Deprecated feature “Control Domain on Solaris 10”

It is highly recommended to use Solaris 11 for Control Domains.

VDCF still supports Solaris 10 Control Domains but the use is deprecated. Support for Solaris 10 will be removed in a future VDCF Release.

Deprecated operations

```
osmon -c show
```

Replaced by `osmon -c show_fs`

Deprecated flags

```
node -c inactivate force
```

The force flag has no effect anymore. A Node can be inactivated anytime. This does only change the Node cState to INACTIVE.

Don't use the above 'force' flag anymore. It will be removed **in the next major** VDCF release.

```
gdom -c addnet ipmp
```

The 'ipmp' flag has no impact anymore. Addnet automatically uses IPMP if the Control Domain or IO Domain Network has IPMP configured.

14 History: What was new in Release 5.6 (27. October 2015)

14.1 Oracle Solaris 11.3 supported

This VDCF Release supports the recently released Solaris 11 Update 3.

14.2 Oracle VM for SPARC / LDom 3.3 supported

This VDCF Release supports the current release of LDom Version 3.3.

14.3 Extended OS Monitoring

With this Release the OS Monitor (osmon) collects state information of SMF services and Datasets (ZPOOLS). Alerting can be enabled if SMF services are not properly running or a Dataset is getting full.

14.4 Support for split IO Guest Domains

If a Solaris SPARC CMT System is setup with Root IO domains VDCF will configure Guest Domains automatically with redundant IO (disks and network) for high availability. Split IO Domains are supported if MPXIO LUNs are used.

You need to enable this feature by setting `IODOM="TRUE"`

14.5 New operations

<code>gdom -c destroy</code>	Completely destroy a guest domain including datasets.
<code>osmon -c show_fs</code>	Used to show OS Monitoring filesystem data (this operation does the same as the former <code>osmon -c show</code> command)
<code>osmon -c show_smf</code>	Used to show OS Monitoring SMF service data
<code>osmon -c show_dataset</code>	Used to show OS Monitoring Dataset (zpool) state data
<code>osmon -c show_server</code>	Used to show OS Monitoring objects for a server
<code>osmon -c summary</code>	Used to show a summarized view on critical objects (vServer filesystems, datasets and SMF services)

14.6 Enhanced operations

<code>gdom -c console</code>	Guest domains can have their own console escape character. See <code>GDOM_ESCAPE_DEFAULT</code> for more details.
<code>gdom -c modify</code>	New autoboot argument to change the <code>ldom</code> auto-boot flag of guest domains
<code>gdom -c show</code>	The disk id from the <code>ldm</code> configuration is displayed for each disk
<code>gdom -c shutdown cdom=</code>	With this Release only non IO domains are shutdown when using the ' <code>cdom</code> ' argument. To also shutdown the IO domains use the additional argument ' <code>iodom</code> '.
<code>ipsadm -c show_repo</code>	New arguments added: use ' <code>local</code> ' to list all IPS repos on the VDCF server and ' <code>oracle</code> ' to list all SRUs in the official Oracle Solaris Support Repository.
<code>node -c console</code>	Physical nodes can have their own console escape character. See <code>NODE_ESCAPE_DEFAULT</code> for more details.
<code>osmon -c update</code>	New optional arguments (<code>dataset,fs,smf</code>) to only update dataset, filesystem or SMF information.
<code>patchadm -c diff</code>	New ' <code>verbose</code> ' flag to list also architecture dependent patches
<code>rcmon -c disable node=</code>	The ' <code>node</code> ' argument accepts now a comma separated list of node names.
<code>rcmon -c enable node=</code>	The ' <code>node</code> ' argument accepts now a comma separated list of node names.
<code>serverconfig -c exec</code>	You can use the new arguments ' <code>user</code> ' or ' <code>root</code> ' to overwrite the default <code>exec</code> user. See <code>CONFIG_EXEC_COMMAND_DEFAULTUSER</code>

14.7 Deprecated operations

<code>osmon -c show</code>	Replaced by <code>osmon -c show_fs</code>
----------------------------	---

14.8 Deprecated flags

<code>gdom -c addnet ipmp</code>	The ' <code>ipmp</code> ' flag has no impact anymore. Addnet automatically uses IPMP if the Control Domain or IO Domain Network have IPMP configured.
----------------------------------	---

14.9 Configuration variables

New settings

CONFIG_EXEC_COMMAND_DEFAULTUSER

For security reasons you may set a non-root user as default user for your remote command execution. Used by `serverconfig -c exec command=`

Example: `export CONFIG_EXEC_COMMAND_DEFAULTUSER="nobody"`

Allowed values: a Solaris user existing on all systems
default: root

GDOM_AUTOBOOT_DEFAULT

Define the auto-boot value for new Guest Domains

Example: `export GDOM_AUTOBOOT_DEFAULT="FALSE"`

Allowed values: TRUE | FALSE
default: TRUE

GDOM_SHOW_ATTR

Define attributes to be shown on gdom lists:

Example: `export GDOM_SHOW_ATTR="BUILD,OS,PATCH_LEVEL"`

Allowed values: MODEL,BUILD,OS,PATCH_LEVEL,MAU
default: MODEL,OS,PATCH_LEVEL

GDOM_SYSATTR_SHOW

Define attributes to be shown in gdom show detail view:

Example: `export GDOM_SYSATTR_SHOW="AUTOBOOT"`

Allowed values: AUTOBOOT or empty
default: AUTOBOOT

GDOM_ESCAPE_DEFAULT

Define the Guest domain console escape character if it's different from your vServers console escape character.

Example: `export GDOM_ESCAPE_DEFAULT="%"`

Allowed values: Valid escape characters
default: <empty> (escape char defined in VIRTUAL_ESCAPE_DEFAULT is used)

IODOM

Enable support for Root IO domains.

Example: `export IODOM="TRUE"`

Allowed values: TRUE | FALSE
default: FALSE

IPS_DEFAULT_REPO

Define your default remote IPS repository to be used in these `ipsadm` commands:
`ipsadm -c show_repo / create_service / create_build`

Example: `export IPS_DEFAULT_REPO="http://myreposerver:8282"`

Allowed values: Valid IPS repository URL
default: <empty>

NODE_ESCAPE_DEFAULT

Define the node console escape character if it's different from your vServers console escape character.

Example: `export NODE_ESCAPE_DEFAULT="%"`

Allowed values: Valid escape characters
default: <empty> (escape char defined in `VIRTUAL_ESCAPE_DEFAULT` is used)

VIRTUAL_SYSATTR_SHOW

Define attributes to be shown in `vserver show` detail view:

Example: `export VIRTUAL_SYSATTR_SHOW="AUTOBOOT,LOCKED"`

Allowed values: AUTOBOOT,LOCKED
default: AUTOBOOT

14.10 Bug fixes

Several small bug fixes and performance enhancements.

15 History: What was new in Release 5.5 (17. April 2015)

15.1 Oracle VM for SPARC / LDom 3.2 supported

This VDCF Release supports the current release of LDom Version 3.2.

15.2 New Cronjob for package analyze (ACTION required)

This Release introduces an additional cronjob for package analyze.

`vpkgadm_nightly` executes a `vpkgadm analyze` on all systems

It is highly recommended to add the following cronjob:

```
0 2 * * * /opt/jomasoft/vdcf/sbin/vpkgadm_nightly >/dev/null 2>&1
```

In previous Releases the `patchadm_nightly` cronjob did this analyze, which is not the case anymore! On a Solaris 11 only environment you can safely disable the `patchadm_nightly` cronjob.

15.3 ZFS Volumes (ZVOL)

VDCF now supports the use of prepared ZFS Volumes. This allows you to virtualize Systems where no access to central storage is available. The ZFS Volumes must be created by the Administrator and registered using the `diskadm register` operation. If you place your vServer or GDom on ZFS Volumes Migration is not supported.

Details can be found in Chapter 9 in the “VDCF Base - Administration Guide”

15.4 Parsable output for some show commands

The show commands of `cdom`, `cpool`, `dataset`, `diskadm`, `gdom`, `node` and `vserver` can now produce machine readable output. This is useful to feed VDCF information into other knowledge management systems like CMDDBs.

15.5 Performance Enhancement / Parallel Execution

`vserver boot/reboot/shutdown` and `vdcfadm update_node` are now per default working on 10 Objects in parallel for faster execution.

15.6 New Cronjob for diskusage statistics

A new cronjob is available to calculate disk usage data for all systems (Nodes and vServers). This command does generate disk usage data and writes it to a history database and optionally to a comma separated report file (add `-F` to disable report files).

We recommend to run this cronjob once a month:

```
0 0 1 * * /opt/jomasoft/vdcf/sbin/diskusage_update -F -q >/dev/null 2>&1
```

15.7 New operations

<code>diskadm -c statistics</code>	Show disk/LUN usage statistics grouped by Node/vServer and storage tier. To get statistics data you have to enable the diskusage cronjob (see 2.6).
<code>gdom -c attach_root_mirror</code>	Attach mirror disk to a root pool of a Solaris 11 GDom.
<code>node -c upgrade</code>	Update a node with its vServers to a newer Solaris 11 SRU.

15.8 Enhanced operations

<code>cdom -c show</code>	New option 'parsable' and 'header' to get machine readable output New optional argument 'cpool' to list only cdoms of a certain cpool.
<code>cpool -c show</code>	New option 'parsable' and 'header' to get machine readable output
<code>dataset -c show</code>	New option 'parsable' and 'header' to get machine readable output. New optional argument 'vserver' to show datasets from a vserver only.
<code>diskadm -c modify</code>	New optional argument 'tier' to define a tier level for a disk/LUN. Storage tier is used by 'diskadm -c statistics' to group the disks.
<code>diskadm -c register</code>	New 'ZVOL' value for the methods argument to register ZFS Volumes.
<code>diskadm -c show</code>	New option 'parsable' and 'header' to get machine readable output. New optional argument 'vserver' and 'tier' to show only LUNs from either the given vserver or a storage tier.
<code>gdom -c install</code>	New option 'console' to open the console right after the installation starts.
<code>gdom -c modify</code>	Allows to change 'ram', 'cpu' and 'cores' of a detached GDom.
<code>gdom -c show candidates/attach/migrate</code>	New option 'noiocheck' to migrate a Guest domain which is also an IO domain to another Control domain. By default IO domains can't be migrated.
<code>gdom -c show</code>	New option 'parsable' and 'header' to get machine readable output
<code>ipsadm create_repo/update_repo</code>	Additional optional argument 'dir' to be used with the current delivery form of Solaris 11 SRUs (directory with zipped SRU files).



<code>node -c show</code>	New option 'parsable' and 'header' to get machine readable output
<code>node -c install</code>	New option 'console' to open the console right after the installation starts.
<code>node -c register</code>	Argument 'build' is now optional.
<code>nodecfg -c modify</code>	New option 'proxy' and 'clear_proxy' to set or remove proxy definition of a node.
<code>patchadm -c diff</code>	New option 'node' and 'all' to show differences between a node and the vServer running on it.
<code>vdcfadm -c update_node</code>	The operation has been enhanced to update multiple nodes in parallel. You can control this feature by the var VDCFADM_UPDATE_BATCH.
<code>vserver -c boot/reboot/shutdown</code>	Enhanced option 'node' accepts now a list of nodes. The operation has been enhanced to allow to start/stop vServer in parallel. You can control this feature by the var VIRTUAL_CONTROL_BATCH.
<code>vserver -c console</code>	Additional optional arguments 'history', 'tail' and 'follow' to display console history. Requires Solaris 11.2.
<code>vserver -c modify</code>	New option 'autoboot' to switch autoboot flag on or off. New option 'locked' to mark a vserver as not migratable.
<code>vserver -c show</code>	New option 'parsable' and 'header' to get machine readable output

15.9 Configuration variables

New settings

NET_MAPPING_SHOW

Enables or disables to show network mappings in vServer or GDom show output.

Example: `export NET_MAPPING_SHOW="TRUE"`

Allowed values: TRUE|FALSE
default: FALSE

NOAUDIT_FOR_USER

For users defined in this variable VDCF produces no audit records.
Used to avoid unnecessary records by management or reporting jobs.

Example: `export NOAUDIT_FOR_USER="splunk report"`

default: <empty>

VIRTUAL_CONTROL_BATCH

vServer boot, reboot and shutdown has been enhanced. The Operation can be run in parallel.
With the variable VIRTUAL_CONTROL_BATCH, you can set the maximum number of vServer to be stopped or started.

Example: `export VIRTUAL_CONTROL_BATCH=5`

Allowed values: Numbers
default: 10

VDCFADM_UPDATE_BATCH

`vdcfadm -c update_node` operates on multiple Nodes in parallel. With the variable VDCFADM_UPDATE_BATCH, you can set the maximum number of Nodes.

Example: `export VDCFADM_UPDATE_BATCH=5`

Allowed values: Numbers
default: 10

15.10 New configuration files

`/var/opt/jomasoft/vdcf/conf/ntp.cfg`

or

`/var/opt/jomasoft/vdcf/conf/<node>_ntp.cfg`

This files can optionally be used as a NTP configuration file templates.
Use this variables as placeholder for your time servers:

`{NTP_TIMESERVER_1}`, `{NTP_TIMESERVER_2}`

Consult Chapter 13 in the “VDCF Base - Administration Guide” for details.

15.11 Bug fixes

Several small bug fixes and performance enhancements.

15.12 Deprecated feature “Control Domain on Solaris 10”

It is highly recommended to use Solaris 11 for Control Domains.

VDCF still supports using Solaris 10 for Control Domains but the use is deprecated. Support for Solaris 10 will be removed in a future VDCF Release.

16 History: What was new in Release 5.4 (27. August 2014)

16.1 Solaris Update 11.2 supported

This VDCF release supports the current release of Solaris 11: Update 2.

16.2 Solaris 11.2 Unified Archives

Nodes and GDoms can now be installed using Unified Archives. Unified Archives are the Archive implementation on Solaris 11. They are similar to the well known Flash Archives used with Solaris 10.

This allows rapid installations, quick cloning or simply a way to have a disaster recovery solution.

To use the Unified Archives the VDCF Management Server must run on Solaris 11.2.

Details can be found in the additional document "VDCF Base - Administration Guide"

16.3 VDCF Proxy

VDCF Proxy is a feature which allows to use VDCF in different networks without setting up more than one VDCF Management Server and opening multiple ports to all VDCF clients on the firewalls. With VDCF Proxies you have just one entry point to and from the additional networks.

Details can be found in the additional document "VDCF – Proxy" Guide

16.4 High Available Guest Domains using Solaris Cluster

VDCF automatically integrates Guest Domains into the Cluster Configuration on Control Domains where Solaris Cluster 4.1 or 4.2 is running. This way the Guest Domains are made high available and will switch to another Cluster Node if the Primary Node fails.

Details can be found in the additional document "VDCF – Sun/Solaris Cluster" Guide

16.5 Enhancements for Customers using Solaris 10 and 11

For Customers running a mixed environment with Solaris 10 and 11 Systems, several new flags are available:

- filter the Objects by Solaris Release (s10 or s11).
- tag server configurations with Solaris Release and Architecture (sparc/i386) to deploy appropriate configurations only at install time.

16.6 Node SWAP ZPOOL Dataset

For easier management of additional SWAP, VDCF now supports SWAP ZPOOL datasets. It was never easier to add more swap, remove swap or mirror swap to another Storage. Just create your node ZPOOL dataset using the `swap` flag.

16.7 ssh key Support for Console/System Controller Access

For System Controllers where a ssh key is configured, the VDCF commands `node` and `console` can now connect to the console without the need to enter the password.

Supported Console Types: ILOM, ILOMx86, XSCF

The private key file needs to be configured as: `CONFIG_CONSOLE_PRIVKEY`

16.8 New operations

`config -c rename` Rename existing base configurations

16.9 Enhanced operations

`config -c add/modify` Additional optional arguments '`os`' and '`platform`' to better separate configurations.

`cdom/node/gdom/vserver -c show` New option to filter by Solaris Version: '`s10`' and '`s11`'

`cdom/node/gdom -c show` On the list view: Added columns for vServer (`#V`) and/or GDom (`#G`) counters to show the amount of virtual objects running on that server.

`dataset -c create` Support for swap ZPOOL on nodes added with option '`swap`'.

`gdom/vserver -c show candidates` New check for networks is added to the candidate evaluation. The output of GDom and vServer candidates is now nearly identical.

By default only possible candidates are listed. To display all nodes with their candidate results use the new option '`full`'.

`gdom/vserver -c show` Network types of the underlying Node/CDom is displayed in verbose mode

`gdom -c show candidates` New optional argument '`cdom`' to check only a specific target CDom.

`gdom -c attach/migrate` Optimized VirtualSwitch discovery allows migration between CDomS with different VSW setups (IPMP ↔ ETH ↔ AGGR).

`ipsadm -c create_build` New argument '`archive`' to define a Build based on a Solaris 11.2 Unified Archive.

`ipsadm -c show_build` New optional argument '`name`' to display details of a Build (i.e. get all nodes using that Build).

On the list view: Added counters for enabled (`#E`) and activated (`#A`) to show the usage of the builds.

`ipsadm -c show_service` New optional argument '`name`' to display details of an AI Install Service (i.e. get all Builds using that service)

`nodecfg -c discover` New optional argument '`proxy`' to specify the proxy to be used to connect to the Node. SAN LUNs are not discovered anymore by default.

`osmon -c show` New option '`root`' to display only information about root filesystems.

`vdcfadm -c show_config` New option '`output`' to only display variables to customize the VDCF output format.

<code>vserver -c modify</code>	New argument 'build' to change the s11 build used for that vServer.
<code>vserver -c remfs</code>	New argument 'dataset' to remove all filesystems defined on that dataset.

16.10 Configuration variables

New settings

CONFIG_CONSOLE_PRIVKEY

To let VDCF use ssh keys to access system controllers, define the file name of the private key here.

Example: `export CONFIG_CONSOLE_PRIVKEY="/root/.ssh/console/id_rsa"`

default: <empty>

CONFIG_DISCOVER_SANDISK

Defines if `nodecfg -c discover` should discover SAN LUNs. Customers using SANBOOT should set this value to TRUE.

Allowed values: TRUE|FALSE
default: FALSE

GDOM_NET_MACS

Number of alternate mac addresses created on ldm vnet definitions (only for Solaris 11). This settings limits the number of possible vServers with exclusive IP-Stack on a Solaris 11 GDom.

default: 5

VIRTUAL_NET_MAPPING_S11

Default Virtual Net Mapping is defined with VIRTUAL_NET_MAPPING.
Format: "VirtualNet:NodeNet VirtualNet:NodeNet"
To define a default for Solaris 11 vServer which is different from the default for Solaris 10 vServer set this variable

Example: `export VIRTUAL_NET_MAPPING_S11="management:PUBL backup:BACK"`

default: <empty>

VIRTUAL_RCADM_MUST_ON_GDOM

If TRUE, the vServer needs to set the values defined in the variables VIRTUAL_RCADM_MUST_SELECT and VIRTUAL_RCADM_MUST.

If vServers on GDoms should not be checked, set this variable to FALSE.

Allowed values: TRUE|FALSE
default: TRUE

VDCF_DENY_ROOT_USER

If set to TRUE, the root user cannot execute VDCF commands anymore. This is also the case for cronjobs!

Allowed values: TRUE|FALSE
default: FALSE

16.11 New configuration files

```
/var/opt/jomasoft/vdcf/conf/proxy.cfg
```

Here you define your proxy servers. Details can be found in the additional document “VDCF – Proxy” Guide

```
#ProxyName, SocksIP, WebIP  
TEST,10.1.1.1,10.1.1.2
```

16.12 Bug fixes

Several small bug fixes and performance enhancements.

16.13 Removed Support for Solaris 11.0

VDCF will not support Solaris 11 GA Release 11.0 anymore. Installations with Solaris 11 must begin with Version 11.1

17 History: What was new in Release 5.3 (14. March 2014)

17.1 Filesystem usage monitoring

The filesystem usage is discovered regularly using a cronjob and displayed using `vserver -c show`.

Standard/Enterprise and HA customers with access to the VDCF Monitoring package additionally receive the `osmon` command to configure alarming, if a defined filesystem limit is reached.

Details can be found in the additional documents “VDCF – vServer Administration Guide” and “VDCF – Monitoring Guide”

17.2 LDom control domain discovery

Discover your existing LDom control domains to deploy additional guest domains using VDCF. Existing guests are untouched by VDCF. I/O domains are discovered, but not used for deployments. This will be supported in a future VDCF release.

17.3 LDom cpu core configuration

Beside the existing possibility to configure vCPUs for guest domains you can now assign cores and max-cores to guest domains. Using max-cores the Oracle Software Licensing requirements for Virtualization using LDoms can be full-filled.

17.4 VDCF data now in /var/share for Solaris 11.1

If upgrading VDCF on Solaris 11.1 the VDCF data content in `/var/opt/jomasoft` will be moved to the `/var/share/jomasoft` directory and linked back to the original path. If doing a fresh installation it will be created directly like this.

17.5 New commands

`dependadm` **New Command in the VDCF vServer Package for all customers**

Command to configure dependencies between vServers, for example to make sure an Application vServer is shutdown before the Database vServer is stopped. More details can be found in the vServer Administration Guide.

```
dependadm -c show [ vserver=<name> ]  
  
dependadm -c add      master=<vServer list>  
                      slave=<vServer list>  
  
dependadm -c remove  master=<vServer list>  
                      slave=<vServer list>  
  
dependadm -c remove  vserver=<vServer list>
```

`osmon` **Command in the VDCF Monitor Package for Standard/Enterprise/HA customers**

Manages Operation System Monitoring. This Release includes filesystem usage monitoring.

```
osmon -c enable  
  
osmon -c disable  
  
osmon -c status  
  
osmon -c show        [ over=<percent> ]  
                      [ summary ]  
  
osmon -c update        all | node=<node name>
```

17.6 New operations

<code>cdom -c discover</code>	Discover preexisting control domains into the VDCF repository
<code>console -c check</code>	Checks the connection to node system controllers and updates the VDCF repository with firmware and serial number information.
<code>vserver -c shrinkfs</code>	Allows shrinking of vServer vxfs filesystems.

17.7 Enhanced operations

<code>cdom -c create/modify</code>	Additional optional argument 'cores' to configure whole-cores.
<code>cdom -c show definitions</code> are displayed.	Information about existing VSW, VCC, VDS and IO Domain
<code>dataset -c create</code>	New 'size' argument for automatic selection of a free disk, for cloud service providers.
<code>gdom -c adddisk</code>	New 'size' argument for automatic selection of a free disk, for cloud service providers.
<code>gdom -c create/modify</code>	Additional optional arguments 'cores' and 'max-cores' to configure whole-cores to guests.
<code>gdom -c revert</code>	Additional optional flags 'res', 'disk', 'net', 'all' to revert only the required object changes.
<code>hwmon -c enable/disable</code>	'node' argument to enable/disable individual nodes for Hardware Monitoring.
<code>vserver -c revert</code>	New flag 'network' to revert network changes (like purging)

17.8 cronjobs – Recommended Change for patchadm_nightly

The cronjob to analyze Solaris 10 Patches is typically configured to run each day in the morning:

```
# JSvdcf-patch cron
0 7 * * * /opt/jomasoft/vdcf/sbin/patchadm_nightly >/dev/null 2>&1
```

Oracle now typically only releases Solaris 10 Patches around the same time they release a Solaris 11 SRU. This is currently in the 3rd week of each month. Because of this change we recommend to run this cronjob only once a month and change the cron configuration to:

```
# JSvdcf-patch cron
0 7 20 * * /opt/jomasoft/vdcf/sbin/patchadm_nightly >/dev/null 2>&1
```

17.9 Configuration variables

New settings

OSMON_FS_INTERVAL

Defines at what time the osmon check cronjob should be executed.

Default: 30 * * * * (every hour at :30)

OSMON_FS_WARNING

If this filesystem usage (in %) is reached osmon will send a WARNING eMail.

default: 80

VSERVER_CHECK_DEPEND

Defines if dependadm dependency definitions should be checked by the boot, shutdown and patch operations.

default: FALSE

17.10 Bug fixes

Several small bug fixes and enhancements.

17.11 Deprecated flags

hwmon -c status **verbose** verbose has no effect anymore

rcmon -c status **verbose** verbose has no effect anymore

Don't use the above 'verbose' flags anymore, they will be removed **in the next major** VDCF release. The output is now the same with and without the verbose flag.



18 History: What was new in Release 5.2 (9. September 2013)

18.1 LDom Updates

Oracle VM for SPARC / LDom 3.1 supported

This VDCF release supports the current release of LDom (Oracle VM for SPARC): 3.1
LDom 3.1 is delivered in Solaris 11 Update 1 SRU10.

18.2 Solaris 11 Automated Installer on Solaris 10

This feature allows you to install Solaris 11 Nodes from a Solaris 10 VDCF Management Server.
It is based on "Oracle Solaris 11 Provisioning Assistant for Oracle Solaris 10 (OPA)", a tool available from Oracle Support for Customers.

Requirements and Restrictions:
VDCF must be installed on Solaris 10 (1/13) Update 11 or later and currently only SPARC deployments are supported using a VDCF management server.

Details can be found in the additional document "VDCF – Automated Installer on Solaris 10".

18.3 Automated VDCF Setup

For easier setup of new VDCF environments a new tool "setup_vdcf" is available.
This tool automates the customize steps, for example adds values to customize.cfg and configures the required Apache web server.

Details can be found in the "VDCF – Installation Guides" / Chapter 2.5.1

18.4 VDCF Installation Summary

At the end of installations of Nodes, Guest Domains and Zones a installation summary is provided on the Console, to let the System Administrator know if the installation steps are executed as expected.

```
VDCF: Installation Summary
=====
VDCF Steps                : 20 of 20 successful
WARN: Custom Steps        : 1 of 7 failed
Failed Script:  setup_iscsi_initiator, Exit Code: 5
=====
```

18.5 New Node configuration attributes “InventoryNumber” and “DataCenter”

Inv-No

The 'Inv-No' is optional and can be used to add your internal Inventory number to the VDCF repository. This attribute is enabled by adding

```
export NODE_INVNO_ENABLED="TRUE" to customize.cfg
```

This InvNo is displayed using nodecfg only, if you add 'INVNO' to the NODECFG_SHOW_ATTR variable in customize.cfg

DataCenter

The 'DataCenter' is optional and can be enabled by creating the file
`/var/opt/jomasoft/vdcf/conf/datacenter.cfg`.
In this file you define the allowed datacenters in your environment.

```
-bash-3.2$ cat /var/opt/jomasoft/vdcf/conf/datacenter.cfg
#NODE DataCenters
#DCName,default -> Default Datacenter
#DCName        -> additional DataCenter
#DCName allowed characters and number, no special characters
ZUERICH,default
NEWYORK
SINGAPORE
SANFRANCISCO
```

This DataCenter is displayed using nodecfg only, if you add 'DATACENTER' to the NODECFG_SHOW_ATTR variable in customize.cfg

18.6 New operations

`rcadm -c convert_pool`

Converts manually created CPU pool configurations (using `poolcfg`) into VDCF supported `rcadm` dedicated-cpu configurations. May be used when importing existing Solaris Zones into VDCF.

18.7 Enhanced operations

`build -c show_bootserver`

Additional optional argument `'boot_server'` to get a list of all builds based on that bootserver

`dataset -c attach`

Additional optional argument `'newname'` to rename the dataset (ZPOOL and VXVM only)

`ipsadm -c create_service`

Based on the arguments `'patchlevel'`, `'platform'` and `'repository'` a service may be created from a repository and avoiding the AI-ISO-File download.

`ipsadm -c show_repo`

Additional optional argument `'ai-pkg'` to list the available Patch levels of the AI packages.
The group packages are not listed by default anymore.
Use the new `'groups'` flag to list them.

`node -c show name=
gdom -c show name=`

Now also lists node datasets

`nodecfg -c add`

Optional argument `'noprofile'` to add a new Node to the VDCF repository without using a Node profile. The Node is added just using the information from the node discover.
Use `nodecfg -c modify_net` to define the network types for your network interfaces.

`nodecfg -c discover`

Additional optional argument `'add'` to directly add a new Node to the VDCF repository without using Node profiles. The Node is added just using the information from the node discover.
Use `nodecfg -c modify_net` to define the network types for your network interfaces.

`nodecfg -c modify`

Additional optional arguments `'invno'` and `'datacenter'`

`rcadm -c commit`

Now adds and removes dedicated-cpus on the fly in running vServers. No reboot required anymore.

`rcadm -c show`

Additional optional arguments `'all-states'` to include all vServers, independent of `VIRTUAL_RCADM_SHOW_IGNORE_CSTATES`

`rcadm -c statistics`

Additional optional arguments `'all-states'` to include all vServers, independent of `VIRTUAL_RCADM_SHOW_IGNORE_CSTATES`

18.8 Configuration variables

Changed defaults

DATASET_DEFAULT_TYPE

old default: DISKSET
new default: ZPOOL

If you still would like to use DISKSET as your default, add
`export DATASET_DEFAULT_TYPE="DISKSET"` to `customize.cfg`

New settings for Solaris 10 and Solaris 11

CONFIG_CONSOLE_USER

Used as default for System Controller/Console user at console -c add.

default: admin

DISKS_DEREGISTER_INVISIBLE

Deregisters all invisible disks at `diskadm -c register`, if set to TRUE.

Allowed values: TRUE|FALSE
default: FALSE

NODE_INVNO_ENABLED

Enables InvNo attribute for `nodecfg -c add/modify`

Allowed values: TRUE|FALSE
default: FALSE

NODECFG_SHOW_ATTR

Defines the attributes printed at `nodecfg -c show`

Allowed values: VENDOR|MODEL|HOSTID|SERIAL|INVNO|DATACENTER|LOCATION
default: MODEL,HOSTID,SERIAL,LOCATION

VIRTUAL_RCADM_SHOW_IGNORE_CSTATES

Defines which vServer states should be ignored at `rcadm -c show/statistics`

Allowed values: ATTACHING|DEFINED|DETACHED|DETACHING
default: DETACHED

18.9 Bug fixes

Several small bug fixes and enhancements.

19 History: What was new in Release 5.1 (28. February 2013)

19.1 Solaris and LDom Updates

Solaris 11.1 supported

This VDCF release supports the current release of Solaris 11: Update 1.

Solaris 10 Update 11 (1/13) supported

This VDCF release supports the current release of Solaris 10: Update 11 (1/13).

Oracle VM for SPARC / LDom 3.0 supported

This VDCF release supports the current release of LDom (Oracle VM for SPARC): 3.0

19.2 New RBAC profile “VDCF readonly”

A new VDCF RBAC profile is delivered. Using this profile a user is allowed to execute the displaying operations. Useful for external review user, new employees and engineers on production environments.

```
-bash-3.2$ profiles | grep readonly  
VDCF readonly
```

19.3 New operations

<code>build -c show_bootserver</code>	Lists the existing boot server directories.
<code>build -c remove_bootserver</code>	Removes an unused boot server directory.
<code>diskadm -c init</code>	Used to initialize new, unused Veritas/DMP disks.
<code>diskadm -c label</code>	Used to label new, unused MPXIO or ISCSI disks.
<code>vserver -c destroy</code>	Completely uninstalls and removes a vServer.

19.4 Enhanced operations

<code>config -c add/modify</code>	Additional optional argument ' <code>comment</code> ' to set a comment
<code>gdom -c console</code>	Additional optional arguments ' <code>history</code> ', ' <code>tail</code> ' and ' <code>follow</code> ' to display console history. Requires LDom 3.0.
<code>ipsadm -c update_repo</code>	Additional optional flag ' <code>all-pkgs</code> ' to sync all packages of all publishers. By default only solaris packages are updated.
<code>nodecfg -c add</code>	Additional optional flag ' <code>setprobes</code> ' to define IPMP probes addresses for customers using <code>CONFIG_IPMP_LINK_BASED_ONLY="TRUE"</code>
<code>nodecfg -c create_profile</code>	Additional optional flag ' <code>setspeed</code> ' to define a fix speed (for example <code>100fdx</code>) for individual network interfaces.
<code>patchadm -c download</code>	Additional optional argument ' <code>id</code> ' to download one Solaris 10 patch using the patch-id.
<code>serverconfig -c add/ modify/create_group/modify_group</code>	Additional optional argument ' <code>comment</code> ' to set a comment
<code>vserver -c attach</code>	Additional optional argument ' <code>zbe</code> ' for Solaris 11 only. Using <code>zbe</code> you specify which Zone Boot Environment to use. By default the last used (active) Zone Boot Environment is selected.
<code>vserver -c commit</code>	Additional optional flag ' <code>console</code> ' used together with the boot flag starts the vServer console after booting the vServer.
<code>zfsadm -c rollback</code>	Additional optional flag ' <code>childs</code> ' to apply the operation to the given and all his child filesystems.

19.5 New configuration variables

Settings for Solaris 10 and Solaris 11

CONFIG_IPMP_LINK_BASED_ONLY

Used at `nodecfg -c add`. Should be set to TRUE, if only link-based IPMP is used. Using this setting `nodecfg -c add` doesn't ask about IPMP test addresses.

Allowed values: TRUE|FALSE
default: FALSE

CONFIG_LINK_SPEED_AUTO_ONLY

If set to TRUE, `nodecfg -c create_profile` doesn't ask about network interface speeds anymore. AUTO speed is always used.

Allowed values: TRUE|FALSE
default: TRUE

CPOOL_CHECK_DISK_PER_LOCATION

If set to TRUE, `cpool -c check` does the disk visibility check based on the node location. This setting is typically required if SAN-Based Mirroring is used.

Allowed values: TRUE|FALSE
default: FALSE

DATASET_CONCAT_CHAR

Character used to combine vServer and dataset name, when creating a new dataset.

Default: "_" (Underscore)

NODE_NET_HOSTNAME_POSTFIX

Used at `nodecfg -c add`. Using this hostname postfixes the ip address is discovered based on the Network Type.

default: "MNGT:-mngt,BACK:-back"

Settings for Solaris 11

NETWORK_AGGR_MODE

Supported for Solaris 11.1 Nodes.
Defines the type for Network Link Aggregations

Allowed values: trunk|dImp
default: trunk

19.6 Bug fixes

Several small bug fixes and enhancements.

20 History: What was new in Release 5.0 (21. June 2012)

20.1 New feature “Solaris 11 support”

If VDCF is installed on a Solaris 11 system, you are able to manage Solaris 10 and Solaris 11 servers. VDCF installed on a Solaris 10 system can only manage Solaris 10 servers.

IPS, AI and build management for Solaris 11

For Solaris 11 Nodes, the build concept of VDCF is based on IPS repositories and the Solaris Automated Installer. With the `ipsadm` command you manage your Solaris 11 repositories, services and builds.

See Chapter 4.2 in the VDCF Base Administration Guide for details.

Solaris 11 node installation

Instead of using the VDCF `flash` command, Solaris 11 installations are prepared using new operations of the `node` command.

See Chapter 6.4.2 in the VDCF Base Administration Guide for details.

20.2 New feature “Solaris package information”

Now VDCF stores package information of all managed systems in the database. With the `vpkgadm` command you can analyze, display, search and compare the package levels.

See Chapter 8 in the VDCF Base Administration Guide for details.

20.3 New feature “vServer import”

Import Solaris zones into VDCF. Imported vServer can be managed the same way as vServers installed by VDCF.

See Chapter 3.9 in the VDCF vServer Administration Guide for details.

20.4 New commands

<code>ipsadm</code>	Command to manipulate IPS repositories, AI install services and Solaris 11 builds.
<code>vpkgadm</code>	Command for Solaris package management.

20.5 New operations

<code>node -c enable_install</code>	Prepare a node for a Solaris 11 (AI) installation. Corresponds to the command: <code>flash -c enable_install</code> for Solaris 10 nodes.
<code>node -c disable_install</code>	Disables a node from Solaris 11 (AI) installation. Corresponds to the command: <code>flash -c disable_install</code> for Solaris 10 nodes.
<code>node -c show_enabled</code>	Show all nodes enabled for Solaris 11 installation. Corresponds to the command: <code>flash -c list_active</code> for Solaris 10 nodes.
<code>vserver -c import</code>	Import Solaris zones into VDCF.

20.6 Enhanced operations

<code>diskadm -c mark</code>	Additional optional flag ' <code>comment</code> ' to set a comment for the marked disk.
<code>vserver -c modify</code>	Additional flag ' <code>group_pkg</code> ' to change a vServers group package. The group package is only used for Solaris 11 installations.

20.7 New configuration variables

Settings for Solaris 11

AI_ADMINUSER

Name of admin user created by AI while installing a Node/vServer.
default: admin

AI_ADMINPWD

Password of admin user. Password must be in shadow format.
default: not set (you have to define one)

VIRTUAL_NETSTACK_DEFAULT_S11

Sets the default network stack for vServer on Solaris 11
Allowed values: SHARED|EXCLUSIVE|PRIVATE
default: SHARED

The corresponding variable for Solaris 10 vServer is called:
VIRTUAL_NETSTACK_DEFAULT

SOL11_SUPPORT_CERT

File path to your Solaris 11 Support certificate (used for IPS repository update)
default: /var/opt/jomasoft/vdcf/conf/Oracle_Solaris_11_Support.certificate.pem

SOL11_SUPPORT_KEY

File path to your Solaris 11 Support key (used for IPS repository update)
default: /var/opt/jomasoft/vdcf/conf/Oracle_Solaris_11_Support.key.pem

Settings for Solaris 10 and Solaris 11

DHCP_CONFIG_MODE

Installing x86 Nodes over the network requires a DHCP configuration.
This variable defines how the DHCP configuration is done.

EXEC (default)	VDCF automatically configures the local DHCP server
SHOW	VDCF displays the required DHCP configuration
NONE	VDCF does not generate any DHCP configuration

20.8 Bug fixes

Several small bug fixes and enhancements.

20.9 Removed operations

The following operations have been removed in VDCF Version 5

<code>vserver -c zfs*</code>	replaced by <code>zfsadm</code> command
<code>build -c list</code> and <code>flash -c list_versions</code>	replaced by <code>build -c show</code>
<code>dataset -c show_lun</code>	replaced by <code>diskadm -c show</code>

21 History: What was new in Release 4.2 (21. December 2011)

21.1 New feature “vPool for Physical Nodes”

Previous releases of VDCF used to support vPool assignments for vServers and GDomains. With this release assignments of Physical Nodes are supported in addition. This feature is disabled by default and needs to be activated by setting the variable `VPOOL_NODE` to “TRUE”

The following commands are subject to the new node authorization checks:
`node, nodecfg, patchadm, cdom, flash, serverconfig`

New operations have been added:

```
vpool -c add_node  
vpool -c remove_node
```

Additional arguments have been provided for existing commands:

```
vpool -c show [ nodes ] [ node= ]  
vpool -c create [ node= ]
```

See Chapter 11.3 in the VDCF Base Administration Guide for details.

21.2 New feature “Serverconfig Execution”

Execute VDCF server configurations on existing systems. With this release the following configuration types are supported: `COMMAND, SCRIPT, FILE, PKG, SERVICES`

This functionality is offered by the new operation '`serverconfig -c exec`'. To use this powerful operation a user requires the VDCF RBAC profile “VDCF serverconfig exec”.

It is highly recommended to enable the vPool feature, to restrict the allowed target systems (vServer, Gdomains, Nodes) to individual users.

See Chapter 5.3 in the VDCF Base Administration Guide for details.

21.3 New feature “patch download using PCA”

By default VDCF uses the Solaris 10 tools '`smpatch`' and '`patchsvr`' to download patches. PCA may be used as an alternative to '`smpatch`' and allows to directly download patches from Oracle even if VDCF is running in a vServer/zone. Set `PATCH_DOWNLOAD_TYPE` to “PCA” to activate this feature.

See Chapter 7.2.3 in the VDCF Base Administration Guide for details.

For additional required firewall rules required for PCA, consult Chapter 4.1 of the VDCF Installation Guide.

21.4 New operations

<code>rcmon -c summary</code>	Display 24h average usage data of nodes and vservers.
<code>serverconfig -c exec</code>	Execute configurations on existing systems.
<code>serverconfig -c show_members</code>	Show all members (servers) of a configuration group
<code>vpool -c add_node/remove_node</code>	Assign Physical Nodes to vPools

21.5 Enhanced operations

<code>diskadm -c register</code>	Additional optional flag ' <code>new</code> ' to register new disks only.
<code>diskadm -c deregister</code>	The <code>deregister all</code> deregisters now all unused disks even when there are other disks still in use on that node.
<code>patchadm -c check</code>	Patches required to successfully migrate vServers are more reliably recognized. Special patches 800xxx are ignored
<code>serverconfig -c remove</code>	The ' <code>name</code> ' flag is now mandatory.
<code>rcadm -c commit</code>	Modifications of the number of CPU's (using the <code>CPUs</code> property) are updated online for running vServers. A vServer reboot is not required anymore.
<code>rcmon -c show</code>	The ' <code>server</code> ' flag is now mandatory.

21.6 Performance enhancements

The following commands have been optimized for better performance:

```
diskadm -c update
diskadm -c register [ new ]
```

22 History: What was new in Release 4.1 (4. November 2011)

22.1 Easier VDCF upgrades using bundles

The VDCF packages can be downloaded as bundles, based on your current VDCF license. (Entry, Standard or Enterprise).

The Upgrade is then fully automated using the `vdcf_upgrade` command. See chapter 3.3 for details.

22.2 Solaris 10 Update 10 (8/11) supported

This VDCF release supports the current release of Solaris 10 : Update 10 (8/11).

22.3 New feature “dataset disk location check”

Additional checks are executed based on the dataset layout and the disk location configuration. One of the checks ensures your mirror datasets are distributed over at least 2 different locations. This feature helps to avoid non-working datasets in disaster scenarios, if one of your data center is unavailable.

See Chapter 8.5. in the VDCF Base Administration Guide for details.

22.4 New workaround for “ZFS data filesystems”

“Due to the Solaris bug **6449301** it is not supported to create a new Solaris Zone with ZFS data filesystems.”

`vserver commit` includes now a workaround for this Solaris bug, to let you define and create the ZFS data filesystem and install the vServer. No manual workarounds are required anymore!

To reinstall an existing vServer including ZFS data filesystems a `vserver commit uninstall` and `vserver commit works`.

22.5 Enhanced operations

`dataset -c show` The dataset location compliance is displayed if disk locations are defined.

23 History: What was new in Release 4.0 (7. April 2011)

23.1 New major feature: High Availability (HA) / Automated Failover

In this release we have implemented a new High Availability feature. It's available as an Enterprise Extension to our customers holding an VDCF Enterprise License.


This feature is controlled by this new command:

hamon (High Availability Monitoring)

Used to check the health of VDCF Nodes. If a Node fails it may initiate evacuation of vServers. The vServers are then distributed to the available Nodes based on their resource usage values (RAM and CPU). This solution fills the gap between manual failover and full featured failover using Cluster Software.

See the 'VDCF Monitoring' Guide and the manpages for more information.

The following picture illustrates the migrations if the M5000 in Data Center1 fails.



23.2 New tools

`vserver_local` Utility for easy failover of VDCF, if you installed VDCF itself in a Zone. More about this feature can be found in the VDCF Installation Guide.

23.3 New operations

`dataset -c update` Updates dataset size information in database with effective values.

`node -c evacuate` Detach all vServers and distributes them to Nodes with free resources. (used also by the hamon feature)

`node -c register` Register a system into VDCF, as alternative to install it using VDCF

`nodecfg -c modify_net` Modify network interface configuration of configured Node

`vserver -c reattach` Attaches multiple vServers to previous Node or control domain

23.4 Enhanced operations

diskadm -c show	To display the disk comments the new optional flag 'comment' may be used. Set the config value <code>DISKS_SHOW_COMMENT="TRUE"</code> to always show the disk comments.
node -c remove	New optional flag 'force' to remove node from patch target configurations.
nodecfg -c modify	Additional attributes to update: 'location', 'serial' and 'hostid'
nodecfg -c discover	new optional argument 'nonroot' to discover a node using user vdcfexec instead of root.
vserver -c addnet	Additional stack type 'exclusive' to create exclusive ip-stack vserver.
vserver -c attach	New optional flag 'boot' to boot vServer after successful attach.
vserver -c boot vserver -c reboot vserver -c shutdown	Additional arguments 'node' and 'cdom' to boot, reboot and shutdown multiple vServer with one command.
vserver -c create	Additional optional argument 'hostid' to define hostid for vServer on Solaris 10 Update 9 or later. Additional optional arguments 'category' and 'priority' for High Availability Monitoring
vserver -c modify	Additional arguments 'hostid' and 'clear_hostid' to set and clear hostid for vServer on Solaris 10 Update 9 or later. Additional arguments 'category' and 'priority' for High Availability Monitoring
vserver -c revert	new argument 'all' to revert all filesystems of a vServer.
zfsadm -c destroy zfsadm -c rollback	Arguments 'filesystem' and 'mountpoint' are now optional. Provide the snapshot name as <code>/zfsds/myfs@mysnapshot</code> to the 'snapshot' argument.