

D-RE Commands Specification

Version E1.02b

2014/02/01

Edited by DEOS R&D Center



DEOS Project

JST-CREST

Research Area

“Dependable Operating Systems for Embedded Systems Aiming at Practical Applications”



Japan Science and Technology Agency

Contents

1. dre-app.....	3
2. dre-sys.....	7
3. dre-cgroup	12

History

Version	変更内容	変更者	日付
1.02	Corrected the omissions and typo	Dependable Embedded OS R&D Center	2013. Sep. 1
1.02a	Corrected the omissions	Dependable Embedded OS R&D Center	2013. Sep. 15
1.02b	Corrected the omissions	Dependable Embedded OS R&D Center	2014. Feb. 1

1. dre-app

NAME

dre-app

SYNOPSIS

dre-app command {command options}

DESCRIPTION

This is the script to operate the Application Containers.

COMMANDS

```
create -n <name> [-b|--base <base_image>] [-veth [-ip|--ip4 <address>]] [-p|--port <port>]
[-c|-g|--cgroup <cgroup>] [-e|--exclude <path> [<mount_point>]] [-e|--exclude <path>
<mount_point>]] ...
```

Create an Application Container as "`<name>`". In addition, the following options are available.

`-b|--base <base_image>`

Specify the Base Image which is used as the base file system of the Application Container as "`<base_image>`". If this option is not used, the Application Container will use the default Base Image.

`-veth|--veth`

Generate a virtual Ethernet device which connects the Host OS and the Application Container. This option enables a different IP Address in the Application Container than that in the Host OS to be set.

`-ip|--ip|--ip4 <address>`

Set the IP Address of the Application Container as "`<address>`". If this option is not used, the IP Address of the Application Container will be assigned dynamically by DHCP. This optional operation needs to use "`-veth`".

`-p|--port <port>`

Set the port number which is used by the Application Container for "ssh" as "`<port>`". If this option is not used, the default value (2200) will be set. This optional operation cannot use "`-veth`".

`-c|-g|--cgroup <cgroup>`

Generate a Control Group in the Application Container named "`<name>`" as "`<cgroup>`".

`-e|--exclude <path> [<mount_point>]`

Mount a "`<path>`" in the Application Container named "`<name>`" at the "`<mount_point>`" in the Host OS. "`<path>`" is not saved/loaded by the function of snapshot. If "`<mount_point>`" is not used, "`<path>`" is mounted to the same path as "`<path>`" in the Host OS.

!!! CAUTION !!!

You must not mount paths in multiple Application Containers to the same path in the Host OS at the same time. If you do, your file system may receive severe damage.

start -n <name>

Start the Application Container named "*<name>*".

console -n <name>

Establish serial communication between the Application Container named "*<name>*" and the Host OS. To use this command, "*<name>*" needs to be running. Press `<Ctrl+a q>` to end the serial communication.

stop -n <name>

Stop the Application Container named "*<name>*".

suspend -n <name>

Suspend the Application Container named "*<name>*" which is running.

resume -n <name>

Resume the Application Container named "*<name>*" which is suspended.

checkpoint -n <name> -t <tag>

Generate a checkpoint of the Application Container named "*<name>*" as "*<tag>*". (Not implemented yet.)

restart -n <name> -t <tag>

Restart the Application Container named "*<name>*" from the checkpoint named "*<tag>*". (Not implemented yet.)

save -n <name> -t <tag>

Generate a snapshot of the file system of the Application Container named "*<name>*" as "*<tag>*".

load -n <name> -t <tag>

Restore the file system of the Application Container named "*<name>*" with the snapshot named "*<tag>*".

remove -b|--base <base_image>

Remove the Base Image named "*<base_image>*" of Application Containers.

remove -n <name> --snapshot -t <tag>

Remove the snapshot named "`<tag>`" of the Application Container named "`<name>`".

commit -n <name> -t <tag> -b <base_image>

Create a new Base Image as "`<base_image>`" from the snapshot named "`<tag>`".

status -n <name>

Get the status of the Application Container named "`<name>`".

getpid -n <name>

Get the PID of the Application Container named "`<name>`".

register -n <name> {-pid|-vpid} <pid>

Register the process in the Application Container named "`<name>`" at "`/proc/daware`". You can specify the PID of the process which is assigned at "`<name>`" or in the Host OS as "`<pid>`". The following options are available to specify "`<pid>`".

- `-pid` Specify a PID which is assigned in the Host OS.
- `-vpid` Specify a PID which is assigned at "`<name>`".

unregister -n <name> {-pid|-vpid} <pid>

Unregister a process in the Application Container named "`<name>`" at "`/proc/daware`". You can specify the PID of the process which is assigned at "`<name>`" or in the Host OS as "`<pid>`". The following options are available to specify "`<pid>`".

- `-pid` Specify a PID which is assigned in the Host OS.
- `-vpid` Specify a PID which is assigned at "`<name>`".

pidlist -n <name>

Get the list of the PIDs in the Application Container named "`<name>`" from "`/proc/daware`".

destroy -n <name>

Remove the Application Container named "`<name>`" from the system. You have to stop "`<name>`" before you remove it. In addition, the following option is available.

- `-c|-g|--cgroup <cgroup>`
Remove the Control Group named "`<cgroup>`" from the Application Container named "`<name>`".

search -n <name>

Check for the existence of the Application Container named "`<name>`".

list -b|--base

List the Base Images of Application Containers in the system.

list [--checkpoint|--snapshot] -n <name>

List the Application Containers in the system. In addition, the following options are available.

{--checkpoint|--snapshot} -n <name>

List the checkpoints/snapshots in the Application Container named "*<name>*".
("--checkpoint" is not implemented yet.)

bind -n <name> -g <cgroup> -pid|--pid <pid>

Bind the process in the Application Container named "*<name>*" whose PID is "*<pid>*" to the Control Group named "*<cgroup>*".

unbind -n <name> -pid|--pid <pid>

Unbind a process which is bound to the Control Group named "*<cgroup>*".

usage [args]

Call "*dre-cgroup usage [args]*". See "3. *dre-cgroup*" for details.

limit [args]

Call "*dre-cgroup limit [args]*". See "3. *dre-cgroup*" for details.

unlimit [args]

Call "*dre-cgroup unlimit [args]*". See "3. *dre-cgroup*" for details.

cgroup -n <name> [args]

Call "*lxc-cgroup -n <name> [args]*". Execute "*man lxc-cgroup*" for details.

info -n <name>

Show information on the settings of the Application Container named "*<name>*".

help

Show help.

2. dre-sys

NAME

dre-sys

SYNOPSIS

```
dre-sys command {command options}
dre-sys command {command options} [-a <app_container>]
```

DESCRIPTION

This is the script to operate the System Containers. In case the option “-a <app_container>” appears, this script calls “dre-app command {command options} -n <app_container>” in the System Container named “<name>”. (See “1. dre-app” for details.)

COMMANDS

```
create -n <name> [-b|--base <base_image>] [-f|rootfs <rootfs>] [-ip|--ipv4 <ip_address>]
[-mask|--netmask <mask>] [-gw|--gateway <address>] [-dns|--dns <address>] [-m|--ram
<memory>] [-if|--if <interface>] [-c|g|--cgroup <cgroup>]
```

Create a System Container as “<name>”. The hostname of the Guest OS is the same as “<name>”. In the initial state, the both of username and password of the Guest OS are “dre”. In addition, the following options are available.

-b|--base *<base_image>*

Specify the KVM image which is used as the base file system of the System Container as “<base_image>”. If this option is not used, the System Container will use the default Base Image.

-f|--rootfs *<rootfs>*

Specify the Partition to be mounted instead of the base image. If “<rootfs>” is specified, the OS installed on the specified partition can be used as the Guest OS. The “-b” option cannot be specified when this option is specified.

The following OS's are verified as the Guest OS.

- Ubuntu 12.04
- CentOS 6.4

If you use “Ubuntu 12.04 Desktop” as the Guest OS, the Display Manager should be “GDM”.

-ip|--ipv4 *<address>*

Set the IP Address of the System Container as “<address>”. If this option is not used, the IP Address of the System Container will be assigned dynamically by DHCP.

-mask|--netmask *<mask>*

Set the subnet mask of the System Container as “<mask>”. This optional operation needs to use with “-ip”. If this option is not used even though it is an “-ip” option, the default value (255.255.255.0) will be set.

-gw|--gateway *<address>*

Set the default gateway of the System Container as “<address>”. This optional operation needs to use with “-ip”. If this option is not used even though it is an “-ip” option, the default value (192.168.1.1) will be set.

-dns|--dns *<address>*

Set the IP address of the DNS server of the System Container as "*<address>*". This option needs to appear with "*-ip*". If this option is not used even though "*-ip*" option appeared, the default value (192.168.1.1) will be set.

-m|--ram *<memory>*

Set the size of the memory allocated by the System Container as "*<memory>*" in megabytes. If this option is not used, the default value (512) will be set.

-if|--if *<interface>*

Set the name of the bridge interface of the System Container. If this option is not used, the default value (br0) will be set.

-c|--cg|--cgroup *<cgroup>*

Generate a Control Group in the System Container named "*<name>*" as "*<cgroup>*".

start -n <name> [-vnc :N]

Start the System Container named "*<name>*". If "*-vnc :N*" is specified, a VNC server starts on display N.

stop -n <name>

Stop the System Container named "*<name>*" which is running. When this command is performed, the Guest OS in "*<name>*" will be shut down.

suspend -n <name>

Suspend the System Container named "*<name>*" which is running.

resume -n <name>

Resume the System Container named "*<name>*" which is suspended.

checkpoint -n <name> -t <tag>

Generate a checkpoint of the System Container named "*<name>*" as "*<tag>*".

restart -n <name> -t <tag>

Restart the System Container named "*<name>*" from the checkpoint named "*<tag>*".

save -n <name> -t <tag>

Generate a snapshot of the System Container named "*<name>*" as "*<tag>*".

load -n <name> -t <tag>

Restore the file system of the System Container named "`<name>`" with the snapshot named "`<tag>`".

remove `{--checkpoint|--snapshot} -n <name> -t <tag> or remove {-b|--base} <base_image>`

Remove the checkpoint/snapshot named "`<tag>`" of the System Container named "`<name>`". Otherwise, remove the Base Image `<base_image>`. If the base image is used by some system containers, it can be removed after removing the system containers.

commit `-n <name> -t <tag> -b <base_image>`

Create a new Base Image as "`<base_image>`" from the snapshot named "`<tag>`".

status `-n <name>`

Get the status of the System Container named "`<name>`".

getpid `-n <name>`

Get the PID of the System Container named "`<name>`".

share `-n <name>`

Share the System Container named "`<name>`" with the NFS server. If "`<name>`" does not exist on the NFS server, upload "`<name>`" in the local host to the NFS server. If "`<name>`" already exists in the NFS server, make it available in the local host. In order to perform migration, both the migration source and the migration destination need to use the same Base Image and to share a system container using the NFS server.

unshare `-n <name>`

Stop the sharing of the System Container named "`<name>`" with the NFS server.

immigrate `-n <name> -p <port>`

Start a virtual machine in the target host of migration in order to perform migration of the System Container named "`<name>`". To enable this function, "`<name>`" needs to be shared with the NFS server in advance. (See "`dre-sys share`" for details.) When the migration is performed, the same environment as "`<name>`" will be generated in the virtual machine. After the migration is completed, the original "`<name>`" is destroyed. To communicate during migration, the port number "`<port>`" is used.

(Note: In order to perform migration, both the migration source and the migration destination need to use the same Base Image.)

migrate `-n <name> -d <target> -p <port>`

Perform the migration of the System Container named "`<name>`". To enable this function, "`<name>`" needs to be shared with the NFS server in advance. (See "`dre-sys share`" for details.) When the migration is performed, the same environment as "`<name>`" will be

generated in the virtual machine running in “<target>”. After the migration is completed, the original “<name>” is destroyed. To communicate during migration, the port number “<port>” is used.

(Note: In order to perform migration, both the migration source and the migration destination need to use the same Base Image.)

kill -n <name>

Forcibly terminate the System Container named “<name>”.

destroy -n <name>

Remove the System Container named “<name>” from the system. You have to stop “<name>” before you remove it. In addition, the following option is available.

-c|-g|--cgroup <cgroup>

Remove the Control Group named “<cgroup>” from the System Container named “<name>”.

search -n <name>

Check for the existence of the System Container named “<name>”.

list [--checkpoint|--snapshot] -n <name>

List the System Containers in the system. In addition, the following options are available.

{--checkpoint|--snapshot} -n <name>

List the checkpoints/snapshots in the System Container named “<name>”.

bind -n <name> -g <cgroup> -pid|--pid <pid>

Bind the process in the System Container named “<name>” whose PID is “<pid>” to the Control Group named “<cgroup>”.

unbind -n <name> -pid|--pid <pid> -g <cgroup>

Release a process which is bound to the Control Group named “<cgroup>”.

usage [args]

Call “dre-cgroup usage [args]”. See “3. dre-cgroup” for details.

limit [args]

Call “dre-cgroup limit [args]”. See “3. dre-cgroup” for details.

unlimit [args]

Call "dre-cgroup unlimit [args]". See "3. dre-cgroup" for details.

info -n <name>

Show the settings of the System Container named "<name>".

help

Show help.

3. dre-cgroup

NAME

dre-cgroup

SYNOPSIS

dre-cgroup command {command options}

DESCRIPTION

This is the script to group any process and manage the resources in groups.

COMMANDS

create -n <cgroup>

Generate a Control Group as "*<cgroup>*".

destroy -n <cgroup>

Destroy the Control Group named "*<cgroup>*".

search -n <cgroup>

Check for the existence of the Control Group named "*<cgroup>*".

list

List the Control Groups on the system.

cgroup -n <cgroup> [args]

Call "*lxc-cgroup -n <cgroup> [args]*". Execute "*man lxc-cgroup*" for details.

bind -n <cgroup> [-pid|--pid <pid>]

Bind the process whose PID is "*<pid>*" to the Control Group named "*<cgroup>*".

usage --cpu|--core|--mem|--memsw [-n <cgroup>] [-pid|--pid <pid>]

Obtain the information of the resource usage. The following resources are available.

--cpu CPU usage (%)
 --core CPU core(s)
 --mem Memory usage (bytes)
 --memsw Memory and Swap usage (bytes)

In addition, the following options are available.

-n *<cgroup>*

Obtain the information of the resource usage of the Control Group named "<cgrou>".

-pid|--pid <pid>

Obtain the information of the resource usage of the process whose PID is "<pid>".

limit --cpu|--core|--mem|--memsw -n <cgrou> <limitation>

Set a limitation on resource usage of the Control Group named "<cgrou>", to be "<limitation>". The following resources are available.

--cpu CPU usage (%): Specify an integer value as "<limitation>".

--core CPU core(s): Specify core numbers using the format of "cpuset.cpus" in the cgroup as "<limitation>".

--mem Memory usage (bytes): Specify an integer value as "<limitation>".

--memsw Memory and Swap usage (bytes): Specify an integer value as "<limitation>".

unlimit --cpu|--core|--mem|--memsw -n <cgrou>

Release a limitation on resource usage of the Control Group named "<cgrou>". The following resources are available.

--cpu CPU usage (%)

--core CPU core(s)

--mem Memory usage (bytes)

--memsw Memory and Swap usage (bytes)

help

Show help.