

# SGE Windows Execution Host Installation

---

## Document revision & Copyright

---

1.0, initial publication, 12. December 2007

1.1, added enable\_windomacc, adjustments for AMD64, 16. January 2008

© by Beat Rubischon [beat.rubischon@dalco.ch](mailto:beat.rubischon@dalco.ch) [mailto:[beat.rubischon@dalco.ch](mailto:beat.rubischon@dalco.ch)] (business),  
[beat@0x1b.ch](mailto:beat@0x1b.ch) [mailto:[beat@0x1b.ch](mailto:beat@0x1b.ch)] (private)

Licensed under the GNU General Public License Version 2.

Most up to date version of this document is located on  
0x1b [<http://www.0x1b.ch/misc/papers/sge/>]

## Introduction

---

The following setup was tested using SGE 6.0u9 and SGE 6.1u2. The Windows software package must be taken out of the “supported” Grid Engine Distribution, but it may be combined with a Queue Master based on the “free” Packages. The Queue Master must run on a Linux (-like) operating system, there is (and will probably never be) a port to Windows.

Two installations were done including their related UNIX environment:

- W2k3R2 SP2, 64 Bit, using SUA
- WXP SP2, 32 Bit, using SFU

SFU is an additional, free of charge download for Windows, but works only for 32 bit Windows 2000, XP and 2k3. SUA is included in W2k3R2, there is an additional component which must be downloaded when installing the SUA environment from the Windows installation media. Bot are based on Interix – Google is your friend in case you need more information or more up to date software packages in your SUA / SFU environment.

Don't expect a perfect working environment. There are too many bugs in Windows, SUA / SFU and the Grid Engine itself to make you happy with this solution. In cases you are free to choose, use a Linux (-like) operating system for your compute nodes.

Some decisions must be done in advance: Do you want to integrate the compute nodes into a Windows Domain or do you want to use with local users? At least SUA does not provide a documented way to change the default domain after installation. While W2k domains worked OK for me, I had some troubles using an NT4 aka Samba domain. See issues at the end of the file.

On the nodes, you need two local users: *Administrator* and the Grid Engine Owner, *sge* in my installation. They are referred as `<hostname>+<user>`, `WINNODE01+Administrator` and `WINNODE01+sge` in my examples. Users out of the primary domain (the one joined during installation of SUA / SFU or the local hostname when no domain was joined) are addressed without leading domain name. So `brubischon` and `CORP+brubischon` are the same in the following examples.

I was not able to share `$SGEROOT` between my Linux Headnode and the Windows based compute nodes. As *Administrator* and *sge* must be local users, they don't have access to network resources – no matter if they are placed on a SMB share using Samba or NFS mounted by SUA. You may have luck in environments lacking a domain, where you are able to serve `$SGEROOT` by Samba with `security = share` and `public = yes`. This means that `$SGEROOT/$SGE_CELL/common` must be kept in sync between the Queue Master and the

nodes. Most important is the file `sgepasswd`.

## Installation

---

### Prepare your domain accounts

- Add home using UNC path to user profile
- optional: add custom shell
  - `net USER brubischon /DOMAIN /USERCOMMENT:"/bin/sh" on Domain Controller`
  - `chsh` on client (did not work)
- Register password for Windows using `regpwd` on all nodes to use `rsh` / `rlogin`
- Register password for SGE using `sgepasswd` on the queue master

### Prepare your Queue Master

- Activate Windows support during installation of Queue Master
- Setup DNS in your cluster, `hosts` does not work for SUA/SFU
- Add `<Windows Hostname>@<Windows Domain>` to `hosts` file
- Create CSP certificates
  - `$SGE_ROOT/util/sgeCA/sge_ca -init -days 365`
  - `$SGE_ROOT/util/sgeCA/sge_ca -user <win_admin_name>`
- Create `/opt/sge/default/common/sge_aliases` for path mapping

```

# tar cvf /root/keys.tar /var/sgeCA
# cd /opt/sge
# tar cvf /root/common.tar default/common/

```

You need these file during the node installation. Copy them to the nodes after SUA / SFU installation.

To activate the the usage of `sgepasswd`, you need to reconfigure the queue master:

```

# qconf -mconf
execd_params          enable_windowacc=true

```

### Installation of W2k3R2

- Installation Windows
- Configure LAN, check DNS(!)
- Join Domain (optional)
- Install Patches
- Install SUA
  - Everything except *Visual Studio Debugger Plugin*
  - Enable `setuid`
  - Enable case sensitive filesystem
- Enable `telnet`, `shell`, `ftp` and `login` in `/etc/inetd.conf` to be able to access this system

- fix regpwd
  - # chmod u-s /bin/regpwd
  - # chmod u+s /bin/regpwd

## Installation of WinXP

- Installation Windows
- Configure LAN, check DNS(!)
- Join Domain
- Install Patches
- Disable Firewall
- Disable DEP ( 820-0697.pdf, Page 126)
  - System Properties, Advanced, Startup and Recovery, Edit, /noexecute=alwaysoff
- Install SFU
  - Custom
    - Utilities
    - Interix GNU Components
      - Interix GNU Utilities
      - Interix GNU SDK
    - NFS
      - Client for NFS
      - Server for NFS
    - Password Synchronization
    - Remote Connectivity
      - Windows Remote Shell Service
    - Authentication tools for NFS
      - User Name Mapping
      - Server for NFS Authentication
      - Server for PCNFS
    - \* Interix SDK
    - Active State Perl
  - Enable setuid
  - Enable case sensitive filesystem
  - Local User Name Mapping Server
    - Password and group files
  - Leave file paths empty
- Enable telnet, shell, ftp and login in /etc/inetd.conf

## Install Grid Engine

- Create local Account sge using lusrmgr.msc, set Password
- Register password using regpwd

```
# cd /
# tar xvf /keys.tar
# mkdir /opt/sge
```

```

# chown WINNODE01+sge /opt/sge
# cd /opt/sge
# gunzip -c /nlge-6_1-common.tar.gz | tar xf -
# gunzip -c /nlge-6_1-bin-windows-x86.tar.gz | tar xf -
# tar xf /changes.tar # optional, see below
# tar xf /common.tar
# mkdir default/spool
# chown WINNODE01+sge default/spool

```

Some changes are needed to run the Execution Host on a 64 Bit Windows using SUA:

```

# tar tvfz changes.tar.gz
lr--r--r--  0 2008-01-11 15:28:55 bin/win32-authenticamd -> win32-x86
lr--r--r--  0 2008-01-11 15:28:55 bin/win32-EM64T -> win32-x86
lr--r--r--  0 2008-01-11 15:28:55 lib/win32-authenticamd -> win32-x86
lr--r--r--  0 2008-01-11 15:28:55 lib/win32-EM64T -> win32-x86
-rwxr-xr-x 11775 2007-06-01 13:54:55 util/arch_variables
-rwxr-xr-x  1818 2007-06-01 14:51:57 util/resources/loadsenors/interix-loadsensor.sh
-rw-r--r-- 29507 2007-06-01 14:09:11 util/install_modules/inst_execd.sh
-rw-r--r-- 94341 2007-06-01 14:05:10 util/install_modules/inst_common.sh
lr--r--r--  0 2008-01-11 15:28:55 utilbin/win32-authenticamd -> win32-x86
lr--r--r--  0 2008-01-11 15:28:55 utilbin/win32-EM64T -> win32-x86

```

Run the installer, there is no template for automatic installation yet:

```

WINNODE01+Administrator@winnode01:/opt/sge# ./install_execd

Welcome to the Grid Engine execution host installation
-----
If you haven't installed the Grid Engine qmaster host yet, you must execute
this step (with >install_qmaster<) prior the execution host installation.
For a sucessfull installation you need a running Grid Engine qmaster. It is
also necessary that this host is an administrative host.
You can verify your current list of administrative hosts with
the command:
# qconf -sh
You can add an administrative host with the command:
# qconf -ah <hostname>
The execution host installation will take approximately 5 minutes.

Checking $$SGE_ROOT directory
-----
The Grid Engine root directory is not set!
Please enter a correct path for SGE_ROOT.
If this directory is not correct (e.g. it may contain an automounter
prefix) enter the correct path to this directory or hit <RETURN>
to use default [/opt/sge] >>
Your $$SGE_ROOT directory: /opt/sge

Grid Engine cells
-----
Please enter cell name which you used for the qmaster
installation or press <RETURN> to use [default] >>
Using cell: >default<

Checking hostname resolving
-----
This hostname is known at qmaster as an administrative host.
Hit <RETURN> to continue >>

Local execd spool directory configuration
-----
During the qmaster installation you've already entered a global
execd spool directory. This is used, if no local spool directory is configured.
Now you can configure a local spool directory for this host.
ATTENTION: The local spool directory doesn't have to be located on a local
drive. It is specific to the <local> host and can be located on network drives,
too. But for performance reasons, spooling to a local drive is recommended.
FOR WINDOWS USER: On Windows systems the local spool directory MUST be set
to a local harddisk directory.
Installing an execd without local spool directory makes the host unuseable.
Local spooling on local harddisk is mandatory for Windows systems.
Do you want to configure a local spool directory
for this host (y/n) [n] >>

Creating local configuration
-----
sge@winnode01.cluster added "winnode01.cluster" to configuration list
Local configuration for host >winnode01.corp.dalco.ch< created.
Hit <RETURN> to continue >>

execd startup script
-----

```

```

We can install the startup script that will
start execd at machine boot (y/n) [y] >>
Installing startup script /etc/rc2.d/S96sgeexecd and /etc/rc2.d/K02sgeexecd

SGE Windows Helper Service Installation
-----
If you're going to run Windows job's using GUI support, you have
to install the Windows Helper Service
Do you want to install the Windows Helper Service? (y/n) [n] >>

Grid Engine execution daemon startup
-----
Starting execution daemon. Please wait ...
starting sge_execd

Adding a queue for this host
-----
We can now add a queue instance for this host:
- it is added to the >allhosts< hostgroup
- the queue provides 2 slot(s) for jobs in all queues
referencing the >allhosts< hostgroup
You do not need to add this host now, but before running jobs on this host
it must be added to at least one queue.
Do you want to add a default queue instance for this host (y/n) [y] >> n

Using Grid Engine
-----
You should now enter the command:
source /opt/sge/default/common/settings.csh
if you are a csh/tcsh user or
# . /opt/sge/default/common/settings.sh
if you are a sh/ksh user.
This will set or expand the following environment variables:
- $SGE_ROOT (always necessary)
- $SGE_CELL (if you are using a cell other than >default<)
- $SGE_QMASTER_PORT (if you haven't added the service >sge_qmaster<)
- $SGE_EXECD_PORT (if you haven't added the service >sge_execd<)
- $PATH/$path (to find the Grid Engine binaries)
- $MANPATH (to access the manual pages)

Grid Engine messages
-----
Grid Engine messages can be found at:
/tmp/qmaster_messages (during qmaster startup)
/tmp/execd_messages (during execution daemon startup)
After startup the daemons log their messages in their spool directories.
Qmaster: /opt/sge/default/spool/qmaster/messages
Exec daemon: <execd_spool_dir>/<hostname>/messages

Grid Engine startup scripts
-----
Grid Engine startup scripts can be found at:
/opt/sge/default/common/sgemaster (qmaster and scheduler)
/opt/sge/default/common/sgeexecd (execd)
Do you want to see previous screen about using Grid Engine again (y/n) [n] >>
Your execution daemon installation is now completed.

```

## Test

You should see the nodes on your master:

```

[root@master ~]# qghost
HOSTNAME          ARCH             NCPU  LOAD  MEMTOT  MEMUSE  SWAPTO  SWAPUS
-----
global            -                -     -     -        -        -        -
master            lx24-amd64       8    0.17  2.0G    367.0M  2.0G    8.7M
node01            lx24-amd64       4     -     2.0G    -        2.0G    -
node02            lx24-amd64       4     -     2.0G    -        2.0G    -
winnode01         win32-x86        1     -    383.4M  -       929.6M  -

```

Start the execution daemon:

```

WINNODE01+Administrator@winnode01:/# /etc/init.d/sgeexecd start

```

After two minutes, the metrics should be visible:

```

[root@master ~]# qghost
HOSTNAME          ARCH             NCPU  LOAD  MEMTOT  MEMUSE  SWAPTO  SWAPUS

```

global	-	-	-	-	-	-	-
master	lx24-amd64	8	0.15	2.0G	364.2M	2.0G	8.7M
node01	lx24-amd64	4	-	2.0G	-	2.0G	-
node02	lx24-amd64	4	-	2.0G	-	2.0G	-
winnode01	win32-EM64T	2	0.01	383.4M	324.7M	929.6M	278.9M

## Submit a job

---

Short example for a job submission script:

```
#!/bin/sh
#$ -N Info
#$ -S /bin/sh
#$ -cwd
#$ -q windows.q
echo -----
echo Current directory
pwd
echo -----
echo Environment
set
echo -----
echo UNIX Binary
ping -c 1 server
echo -----
echo Windows Binary
cd /dev/fs/C/WINDOWS
ping.exe -n 1 server
```

## Open Issues

---

- qrsh not working, will change with SGE 6.2
- Automatic install?

sge\_windows\_execution\_host\_installation.txt · Zuletzt geändert: 16.01.2008 18:32 von 127.0.0.1