



Grid-Ireland test facilities

Stephen Childs

Dept. of Computer Science
Trinity College Dublin



Grid-Ireland TestGrid


- Isolated network for testing
 - Replica sites (public addresses)
 - Experimental sites (private addresses)
- Allows us to test **exactly the same configuration** as real sites will use
 - Invaluable for certification of upgrades
 - Decreases time of real deployment
- Mix of Quattor (replicas) and YAIM (experimental sites)
- Could be useful as a realistic yet controlled secondary test environment for SA3?



Aim of TestGrid project




Replica national infrastructure for certification of Grid MW



WebCom-G
UCC TCD NUIG

Development testbed



eGee
Enabling Grids for E-science

Porting testbed

Isolated from production infrastructure

Replicates essential features of infrastructure

Easily accessible from developers' desktops



TestGrid approach

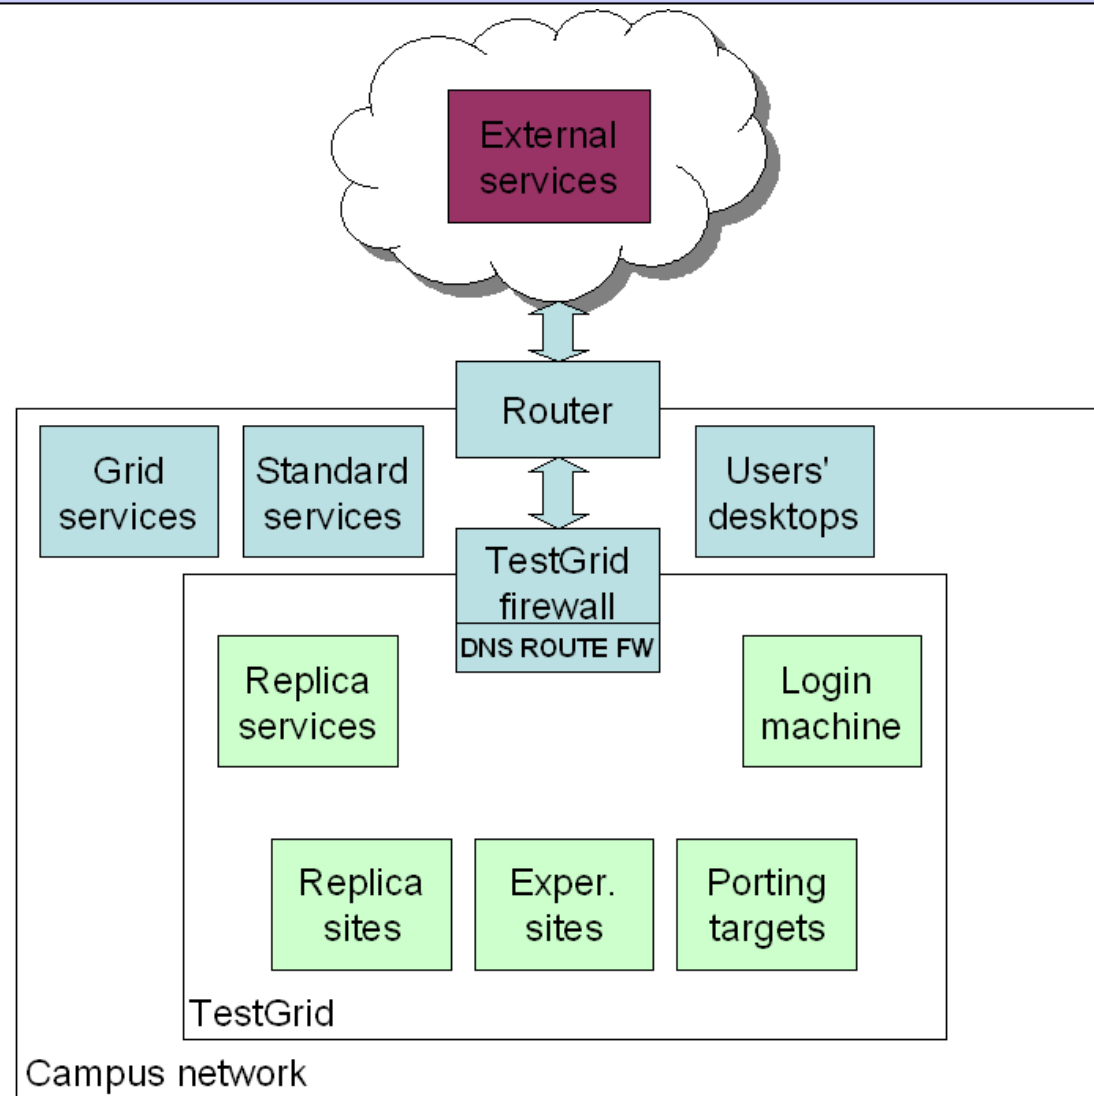
- Single isolated network containing:
 - Replicas of external sites
 - Replicas of Grid infrastructure services
 - Experimental sites
 - Porting build machines and targets
- Providing:
 - Replicas of Grid-Ireland central services
 - Full access to external services
 - Flexibility to implement new services
 - Convenient access from user workstations
 - Same look and feel as real Grid
 - Use of same configuration as real machines

Real addresses

Private addresses



TestGrid structure





TestGrid contents

National Servers (8)

TCD Site (4)

QUB Site (4)

RCSI Site (1)

DIT Site (1)

Virtual site (1)

GT4 site (1)

WN pool (4)

AIX

IRIX

MacOS X

EM64T

Replica central services

134.226.53.0/24

Replica Grid-Ireland sites

134.226.53.0/24, 193.1.229.0/27,
136.206.111.0/24, 147.252.15.0/24, ...

Experimental sites

192.168.0.0/16

Virtual clusters

192.168.0.0/16

Porting targets

192.168.0.0/16



Key technologies

- **Network routing and firewalling**
 - Create isolated network
 - Replicate real network configuration
 - Provide public alias for login machine
 - Masquerade for non-replicated services
- **Virtual machines**
 - Allow replication of large-scale system
 - Allow easy snapshots of system state
- **Fabric management (Quattor)**
 - Perl-like declarative language describes nodes
 - Compiled to XML profile encapsulating OS and MW config
 - Components implement configuration on machine
 - Profiles for all sites stored in one repository
- **Transactional deployment**
 - Deploy version-controlled config. to multiple sites





TestGrid networking

- Routing
 - Multi-homed host; routes for replicas on internal interface
- Firewall
 - Masquerading allows access out to external servers
- Multiple aliases on firewall interface
 - Allows impersonation of site gateways, DNS
- MAC-level translation
 - Implements public aliases to login machines
- DNS server
 - Caches mappings for replica hosts
 - Provides mappings for internal .testgrid domain

Originally configured by hand for each replica network
but error-prone so ...



Quattor router configuration

```
define variable REPLICAS_SITES = {
    sites = nlist();
    sites = npush (
        "DCU",
        nlist("network", "136.206.111.0",
              "gateway", "136.206.111.254",
              "netmask", "255.255.255.0",
              "broadcast", "136.206.111.255"),
        "NUIG",
        nlist("network", "143.203.0.0",
              "gateway", "143.203.7.1",
              "netmask", "255.255.0.0",
              "broadcast", "143.203.255.255")
    );
};

"/system/network/interfaces/eth1/aliases"=
{
    list = self;
    sites = REPLICAS_SITES;
    ok = first(sites,k,v);
    while(ok) {
        list[k]=nlist("ip",sites[k]["gateway"],
                    "netmask",sites[k]["netmask"],
                    "broadcast",sites[k]["broadcast"]);
        ok = next(sites,k,v);
    };
    return(list);
};

"/system/network/interfaces/eth1/route" = {
    sites = REPLICAS_SITES;
    ok = first(sites,k,v);
    while(ok) {
        push(nlist("address",sites[k]["network"],
                  "netmask",sites[k]["netmask"]));
        ok = next(sites,k,v);
    };
};
```

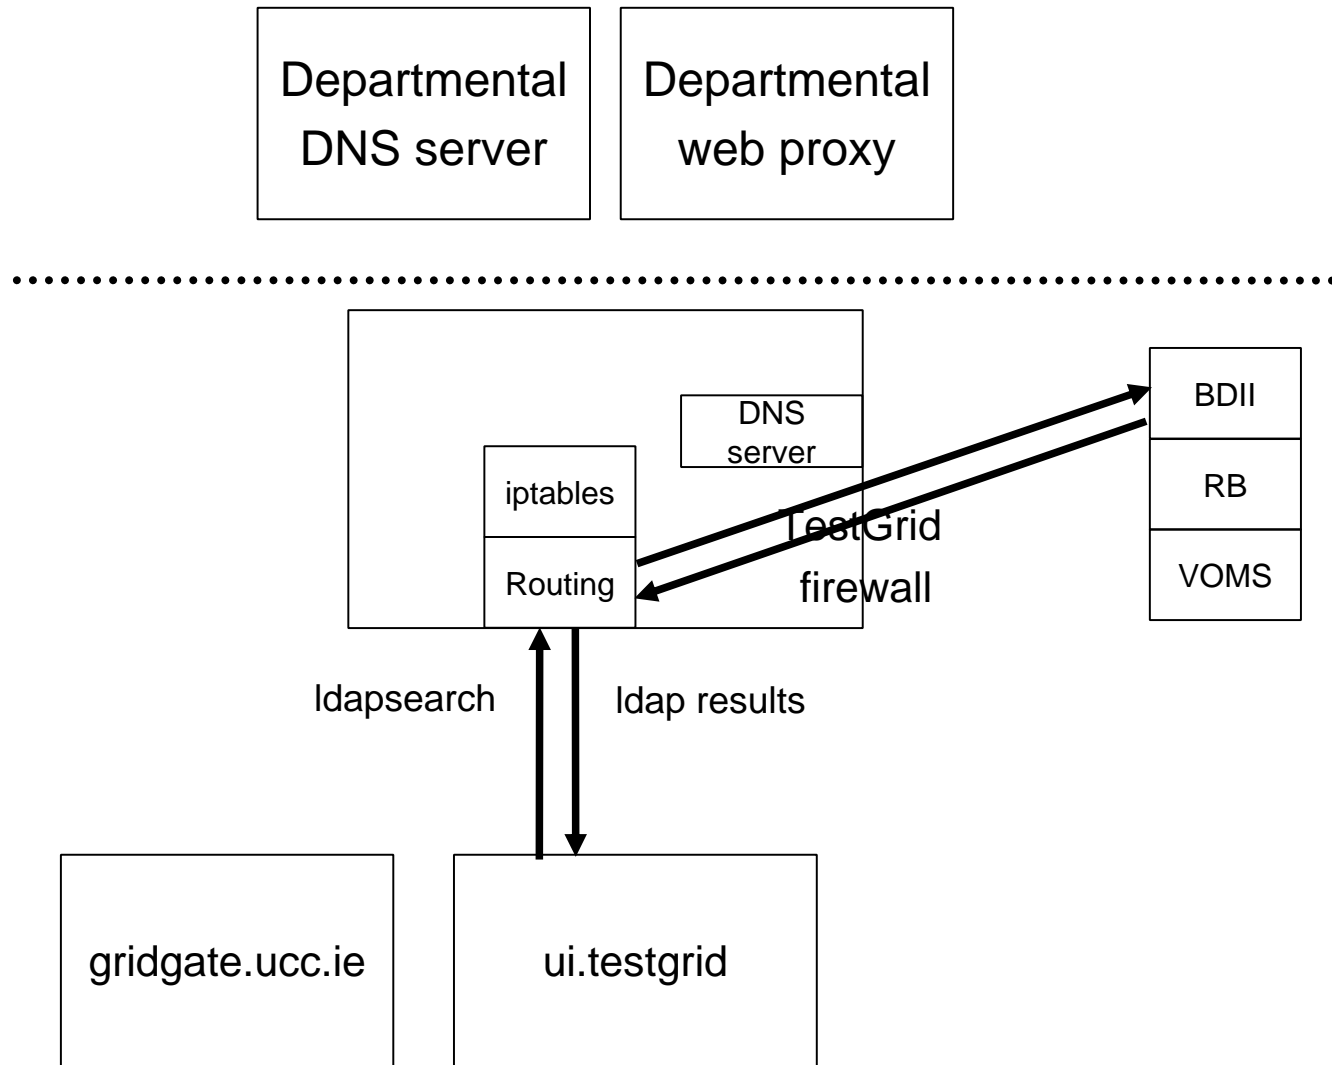
Define network details once

Use to set up aliases

Use to set up routes

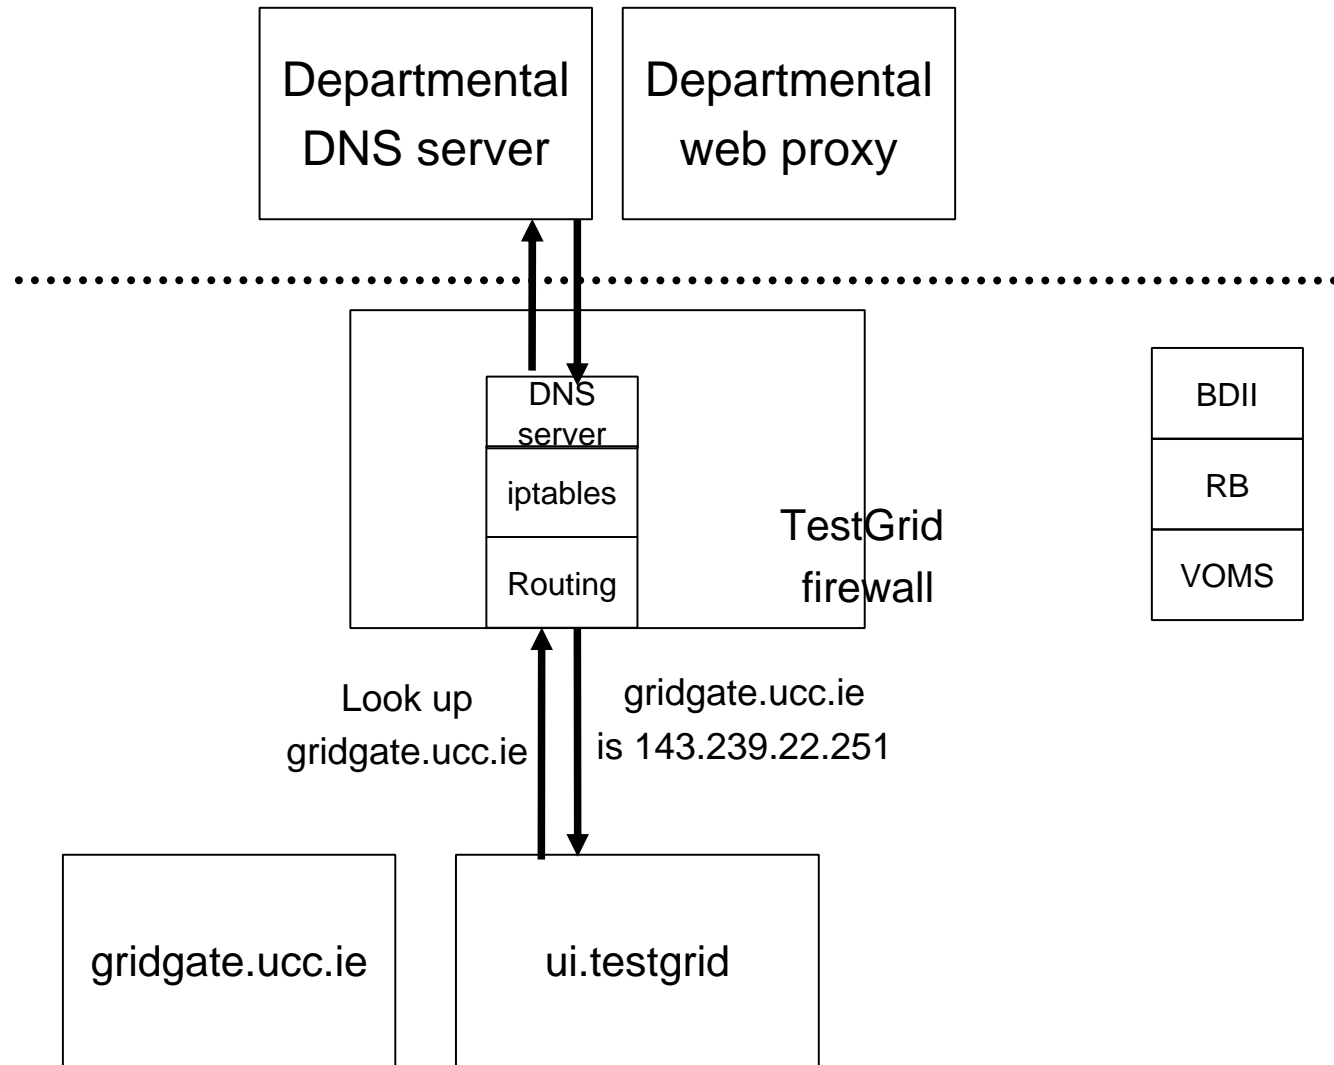


Accessing replicated services



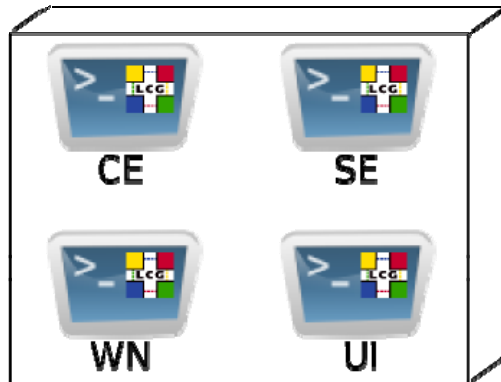


Accessing external services







Applications of TestGrid




Virtual gateway



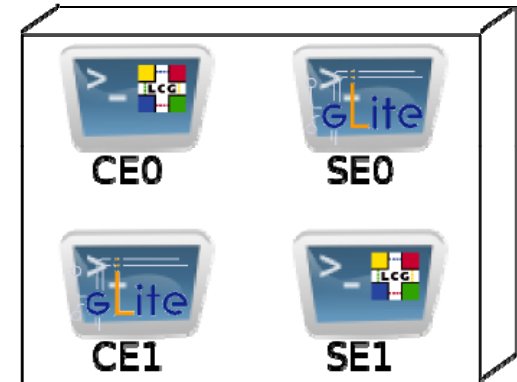
Virtual cluster



Virtual infrastructure



Porting



Middleware testing



Use case 1: Infra. upgrade

- Grid MW has a regular upgrade cycle
- Buggy MW can disrupt production use
- Need full validation before deployment

Release procedure:

1. Edit profiles and upgrade TestGrid
2. Test new release and fix until correct
3. Deploy **same** config. on real Grid

Upgrade of 18 sites from LCG 2.4 to 2.6 took
< 1 week after extensive TestGrid work



Use case 2: Porting

- Grid-Ireland primarily interested in porting worker node (client) software
- TestGrid ideal for testing ports
 - Easily create new VMs
 - Validate against standing infrastructure

Procedure:

1. Create new VMs (build and target) with desired distribution
2. Port software
3. Connect WN to site and run test suite



Management

- Consistency of replicas
 - Quattor configuration profiles: 95%
 - Service-specific synch. (e.g. DBs): 5%
 - Could use VM snapshotting to tighten this
- Server hosts
 - Migrating to Quattor control
 - Developing web tool for managing VMs



Lessons learned & further work

- Even testbed users expect stable infrastructure
 - Separate certification testbed from SW development platform
- Fabric management essential
- Manual provisioning of new VMs problematic

- Add network performance emulation to allow application characterisation
- Better management of virtual machines: GridBuilder web UI



Evaluation

| Requirement | Provided by |
|--------------------|--------------------------------|
| Comprehensiveness | Virtual machines |
| Consistency | Fab. management, TX deployment |
| Isolation | Network configuration |
| Security | Limited access paths |
| Cost-effectiveness | Virtual machines |
| Performance | Xen VM technology |
| Ease of use | Replicating services |
| Flexibility | Experimental sites |



Conclusion

- TestGrid in use for ~ 18 months
- Invaluable for administrators and developers
- Has facilitated rollout of 12 new sites and 3 nationwide upgrades
- Combination of “replica” and “experimental” domains allows variety of applications
- Concept now being reused to create dedicated testbeds for e-learning and certification



Contents of sample VM hosts

Porting

| tg11.testgrid | | |
|---------------|----|----------|
| Name | Id | Mem (MB) |
| Domain-0 | 0 | 251 |
| centos4-build | 1 | 255 |
| fc4-build | 6 | 255 |
| rh73-build | 3 | 255 |
| rh9-build | 2 | 255 |
| sl3-build | 5 | 255 |
| suse9-build | 7 | 255 |

| tg4.testgrid | | |
|--------------|----|----------|
| Name | Id | Mem (MB) |
| Domain-0 | 0 | 251 |
| c01 | 25 | 199 |
| c02 | 27 | 199 |
| c03 | 29 | 199 |
| c04 | 31 | 199 |
| c05 | 33 | 199 |
| c06 | 35 | 199 |
| infogrid2 | 49 | 255 |

Cluster

Virtual site

| tg15.testgrid | | |
|---------------|----|----------|
| Name | Id | Mem (MB) |
| gridgate05 | 1 | 199 |
| gridinstall05 | 21 | 199 |
| gridmon05 | 23 | 199 |
| gridstore05 | 28 | 199 |
| gridwn050 | 24 | 199 |

| tg18.testgrid | | |
|---------------|----|----------|
| Name | Id | Mem (MB) |
| Domain-0 | 0 | 190 |
| infogrid1 | 7 | 255 |
| infogrid3 | 21 | 255 |
| infogrid4 | 13 | 255 |

MW development



Virtual machines

- Extensive use of VMs on real infrastructure and within TestGrid
- Mainly manual configuration, now testing web-based tool (GridBuilder)
- GridBuilder provides
 - Automatic creation of VMs from template images
 - Fast startup & efficient disk usage (COW)
 - Integration with Grid MW config: LCFG, Quattor, (YAIM?)
 - Automatic configuration of VM networking



Grid Ireland Testgrid UI - Mozilla Firefox

https://localhost:4443/cgi-bin/manage.py

Grid Ireland TestGrid cagnode88.cs.tcd.ie

[Home](#)
[Manage](#)
[Create](#)
[Status](#)
[Info](#)

Manage

Running Domains:

| Name | Id | Mem(MB) | CPU | State | Time(s) | Console |
|---------------|----|---------|-----|-------|---------|---------|
| Domain-0 | 0 | 195 | 0 | r--- | 604.9 | |
| gridui_dit_ie | 2 | 127 | 3 | -b--- | 244.0 | 9602 |

Domains in DB:

| nodename | nodetype | baseimage | cowimage | saved | site | owner |
|-----------------------|-----------|----------------------|--------------------------------|-------|----------------|---|
| gridstore_grid_ul_ie | gridstore | /dev/xenvg/gridstore | gridstore_gridstore_grid_ul_ie | no | grid.ul.ie | /C=IE/O=Grid-Ireland/OU=cs.tcd.ie/L=RA-TCD/CN=Stephen O. Childs |
| gridmon_8-tallaght_ie | gridmon | /dev/xenvg/gridmon | gridmon_gridmon_8-tallaght_ie | no | it-tallaght.ie | /C=IE/O=Grid-Ireland/OU=cs.tcd.ie/L=RA-TCD/CN=Stephen O. Childs |
| gridui_dit_ie | gridui | /dev/xenvg/gridui | gridui_gridui_dit_ie | no | dit.ie | /C=IE/O=Grid-Ireland/OU=cs.tcd.ie/L=RA-TCD/CN=Jason McCandless |
| cagnode94 | tylinux | /dev/xenvg/tylbase | tylinux_cagnode94 | no | nosite | /C=IE/O=Grid-Ireland/OU=cs.tcd.ie/L=RA-TCD/CN=Jason McCandless |

Select a stopped domain to restart it:

Memory: MB

Select a domain to remove a filesystem associated with it:

Your name: Jason McCandless
Your cert: /C=IE/O=Grid-Ireland/OU=cs.tcd.ie/L=RA-TCD/CN=Jason McCandless

Done localhost:4443

Grid Ireland Testgrid UI - Mozilla Firefox

Thu Sep 22, 16:18



Grid Ireland TestGrid cagnode88.cs.tcd.ie

[Home](#)
[Manage](#)
[Create](#)
[Status](#)
[Info](#)

Manage

Running Domains:

| Name | Id | Mem(MB) | CPU | State | Time(s) | Console |
|---------------|----|---------|-----|-------|---------|---------|
| Domain-0 | 0 | 195 | 0 | r--- | 604.9 | |
| gridui_dit_ie | 2 | 127 | 3 | -b--- | 244.0 | 9602 |

Domains in DB:

| nodename | nodetype | baseimage | cowimage | saved | site | owner |
|-----------------------|-----------|----------------------|--------------------------------|-------|---------------|---|
| gridstore_grid_ul_ie | gridstore | /dev/xenvg/gridstore | gridstore_gridstore_grid_ul_ie | no | grid.ul.ie | /C=IE/O=Grid-Ireland/OU=cs.tcd.ie/L=RA-TCD/CN=Stephen O. Childs |
| gridmon_8-tallaght_ie | gridmon | /dev/xenvg/gridmon | gridmon_gridmon_8-tallaght_ie | no | 8-tallaght.ie | /C=IE/O=Grid-Ireland/OU=cs.tcd.ie/L=RA-TCD/CN=Stephen O. Childs |
| gridui_dit_ie | gridui | /dev/xenvg/gridui | gridui_gridui_dit_ie | no | dit.ie | /C=IE/O=Grid-Ireland/OU=cs.tcd.ie/L=RA-TCD/CN=Jason McCandless |
| cagnode94 | tylinux | /dev/xenvg/tylbase | tylinux_cagnode94 | no | nosite | /C=IE/O=Grid-Ireland/OU=cs.tcd.ie/L=RA-TCD/CN=Jason McCandless |

Select a stopped domain to restart it:

Memory: MB

Select a domain to remove a filesystem associated with it:

Your name: Jason McCandless
Your cert: /C=IE/O=Grid-Ireland/OU=cs.tcd.ie/L=RA-TCD/CN=Jason McCandless

Done localhost:4443

Grid Ireland Testgrid UI - Mozilla Firefox Thu Sep 22, 16:18 18°