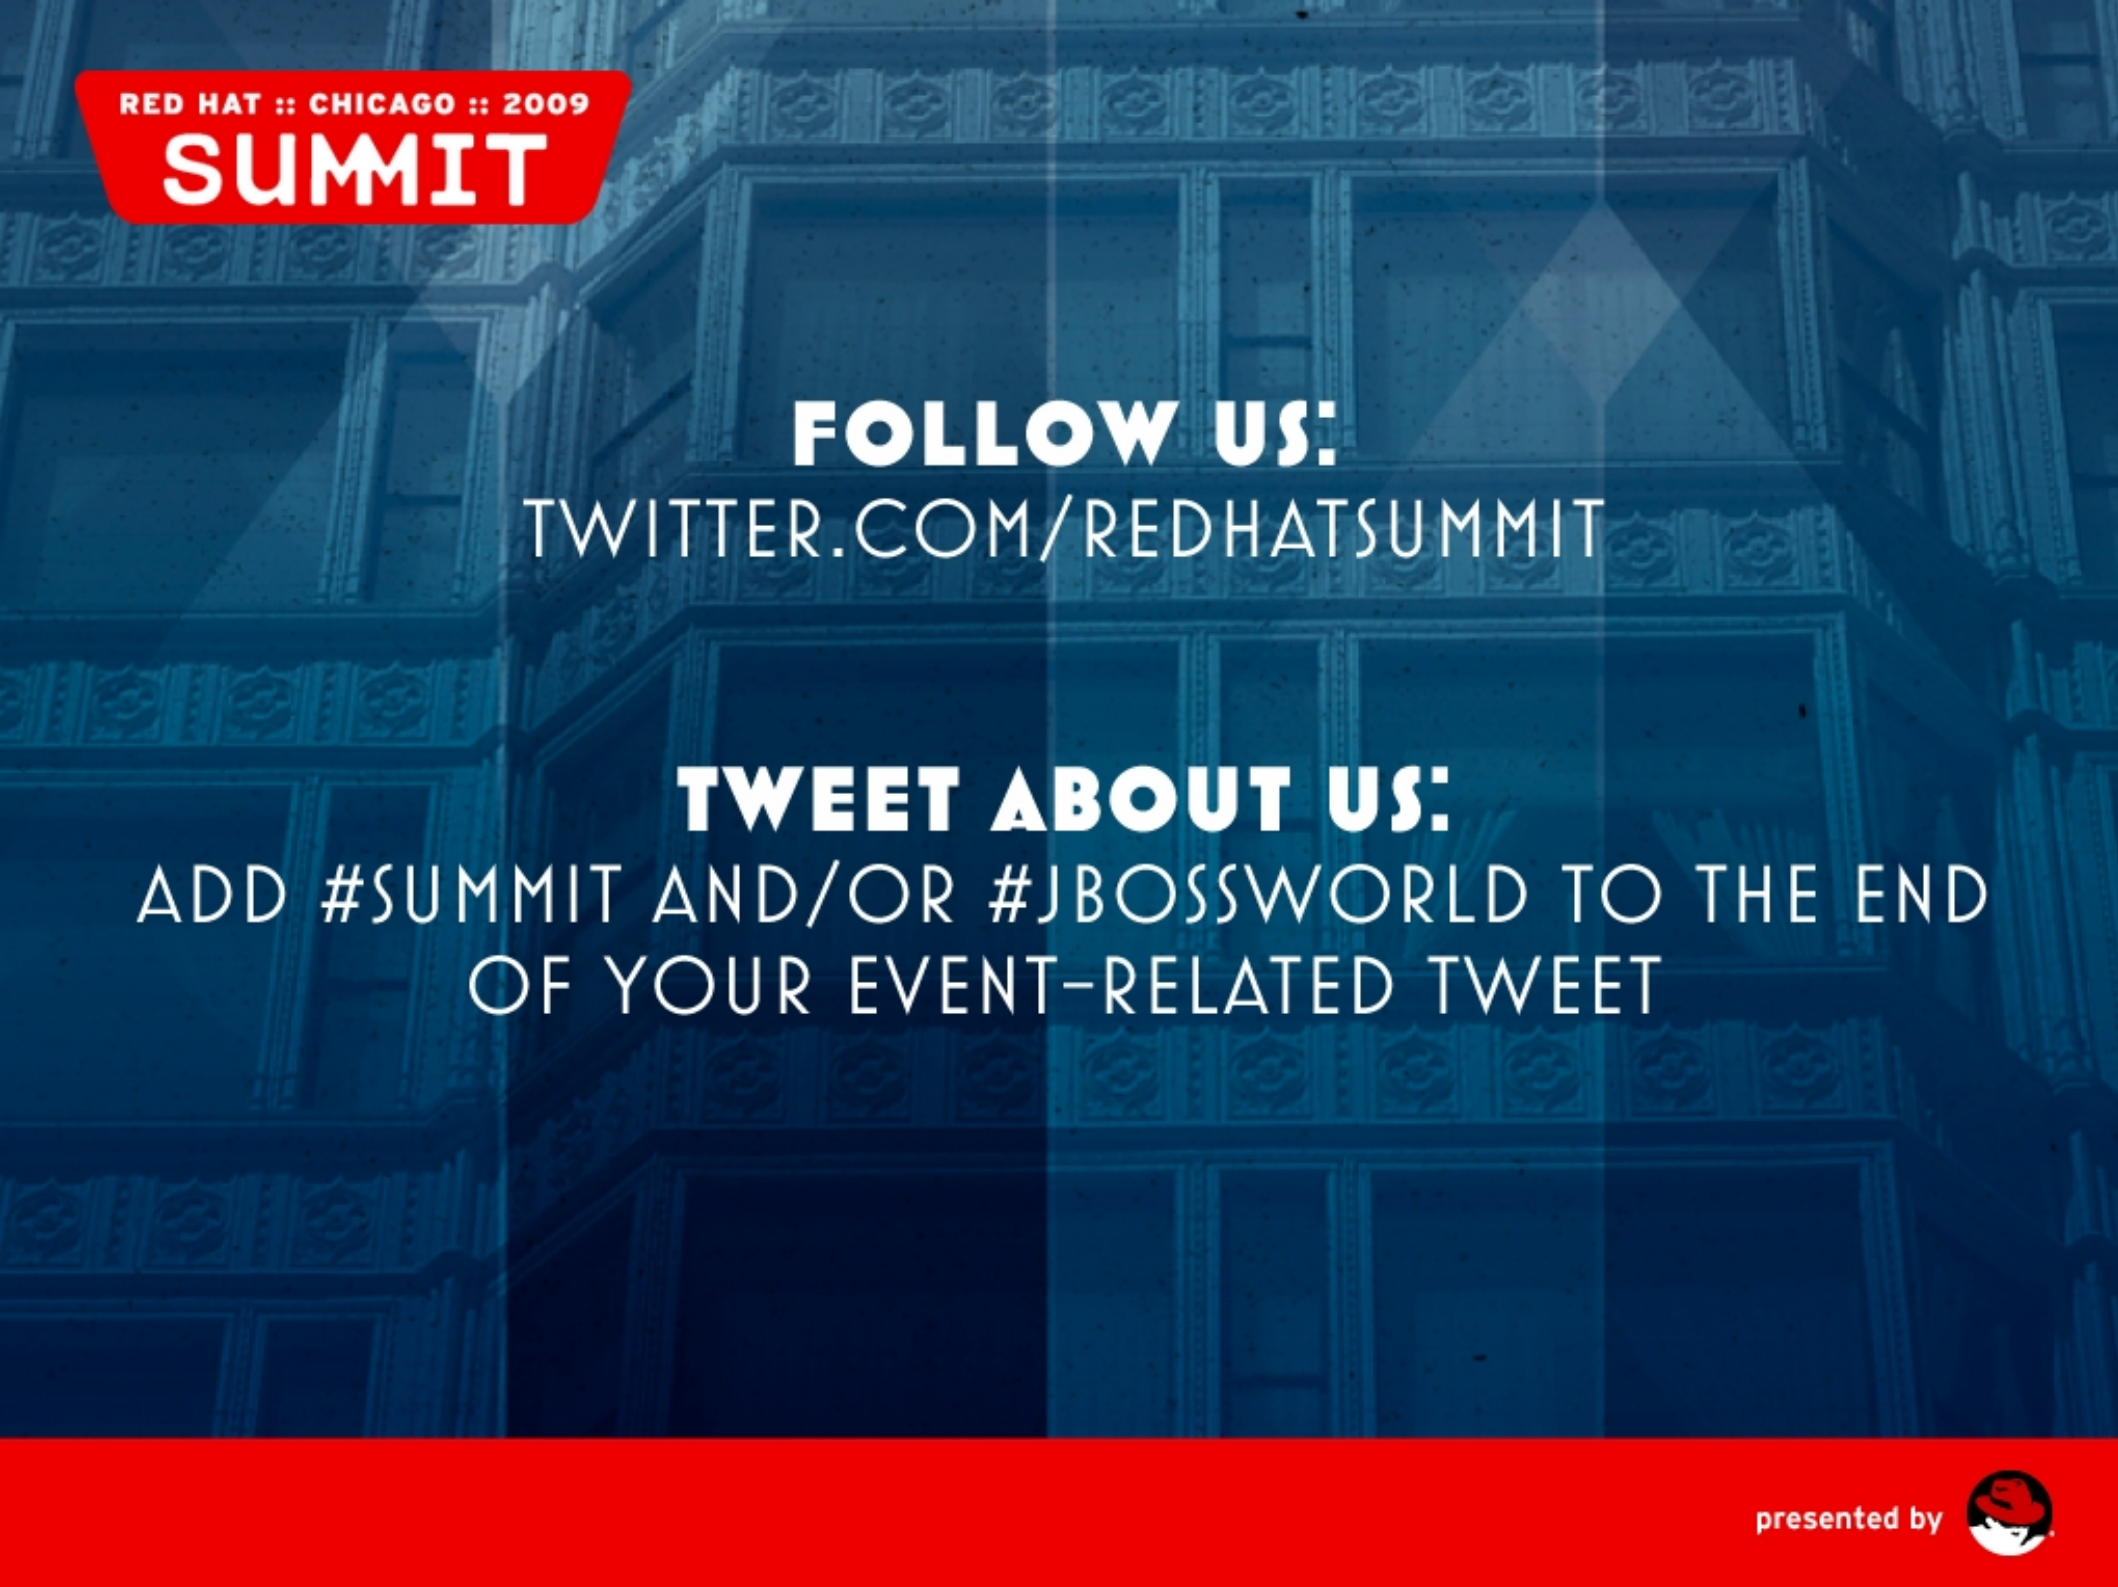
RED HAT :: CHICAGO :: 2009
SUMMIT
FOLLOW US:
TWITTER.COM/REDHATSUMMIT
TWEET ABOUT US:
ADD #SUMMIT AND/OR #JBOSSWORLD TO THE END
OF YOUR EVENT-RELATED TWEET
presented by

RED HAT :: CHICAGO :: 2009
SUMMIT

User-Driven Grid Provisioning

Erich Morisse
Senior Solutions Architect, Red Hat

William Henry
Office of CTO for MRG

Agenda

Complexity in the Data Center

Too Much Efficiency

Reducing Barriers to Change

Complexity in the Data Center

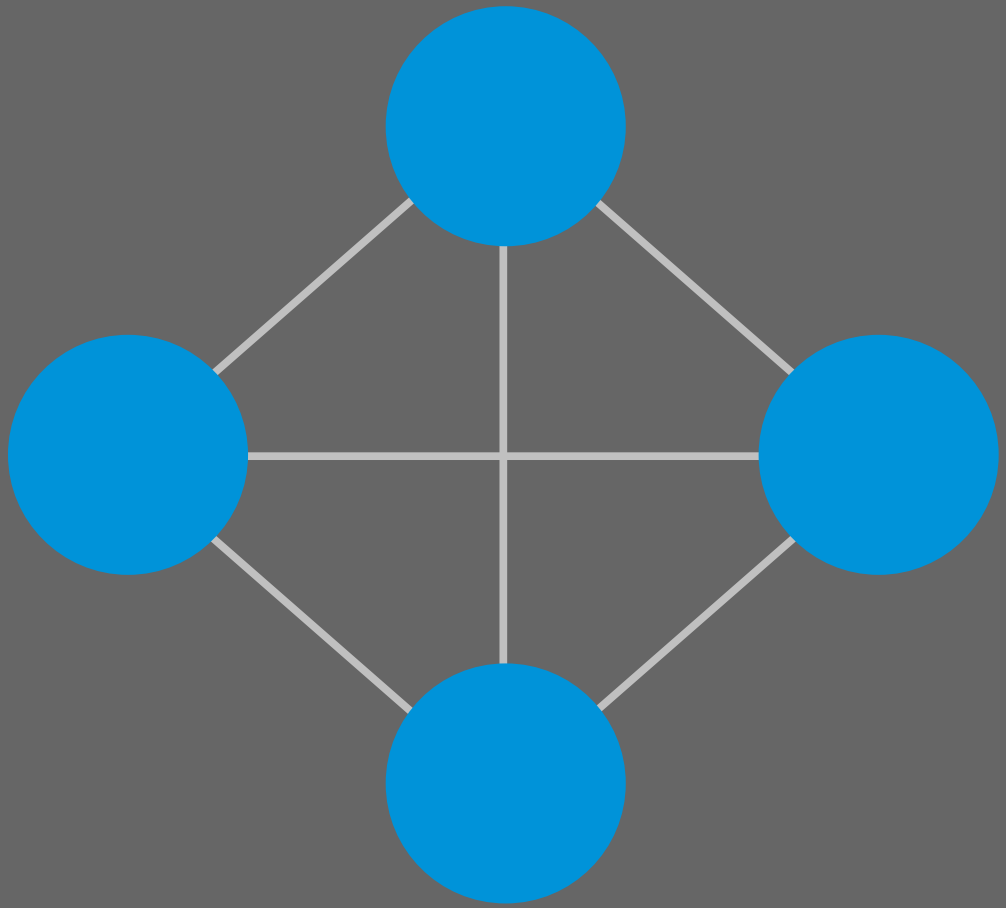

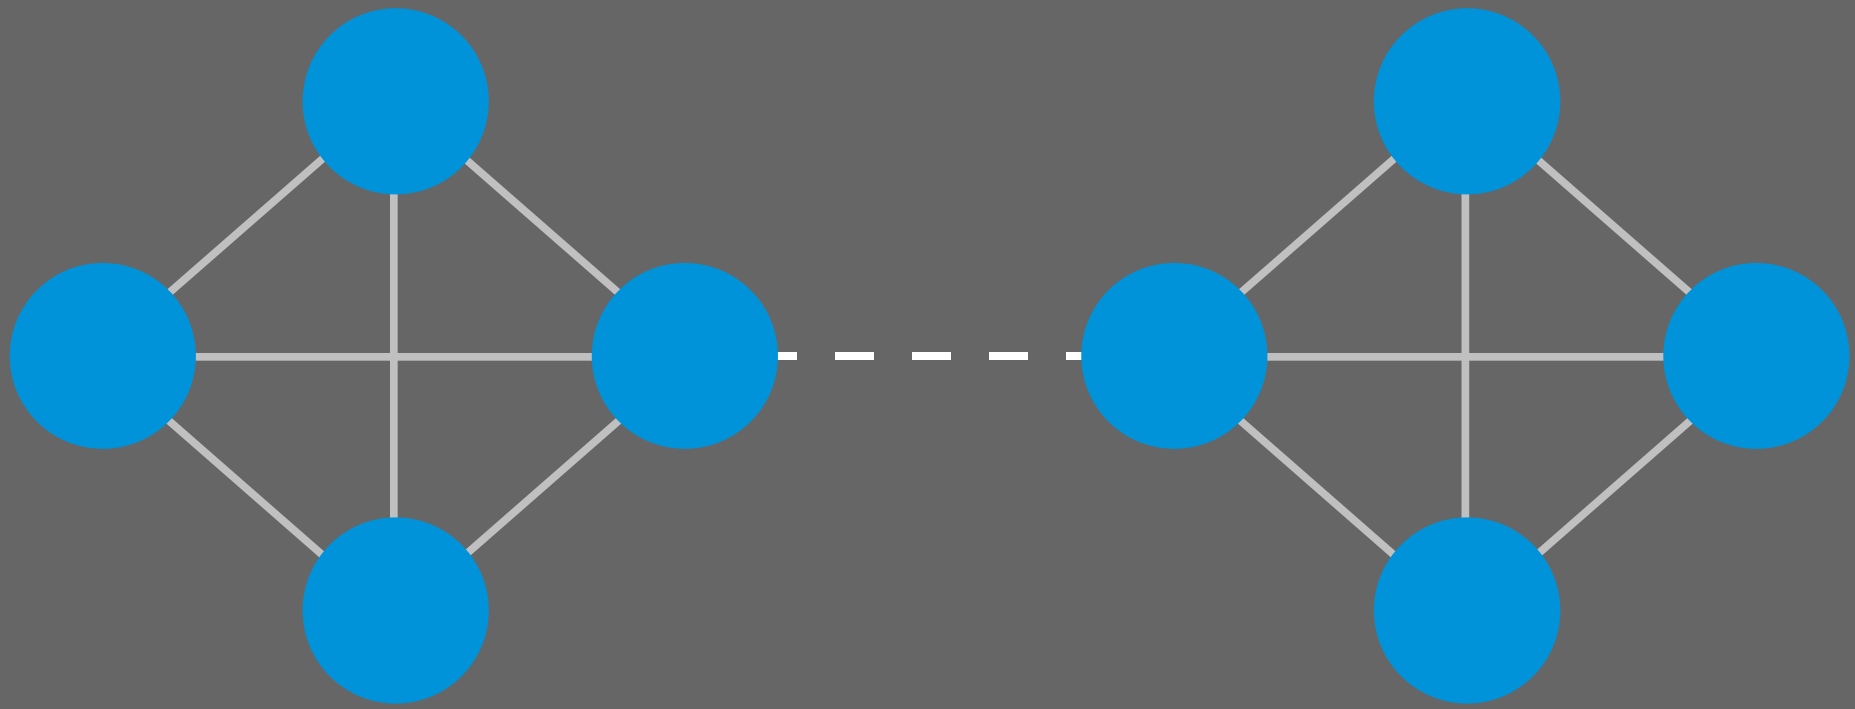

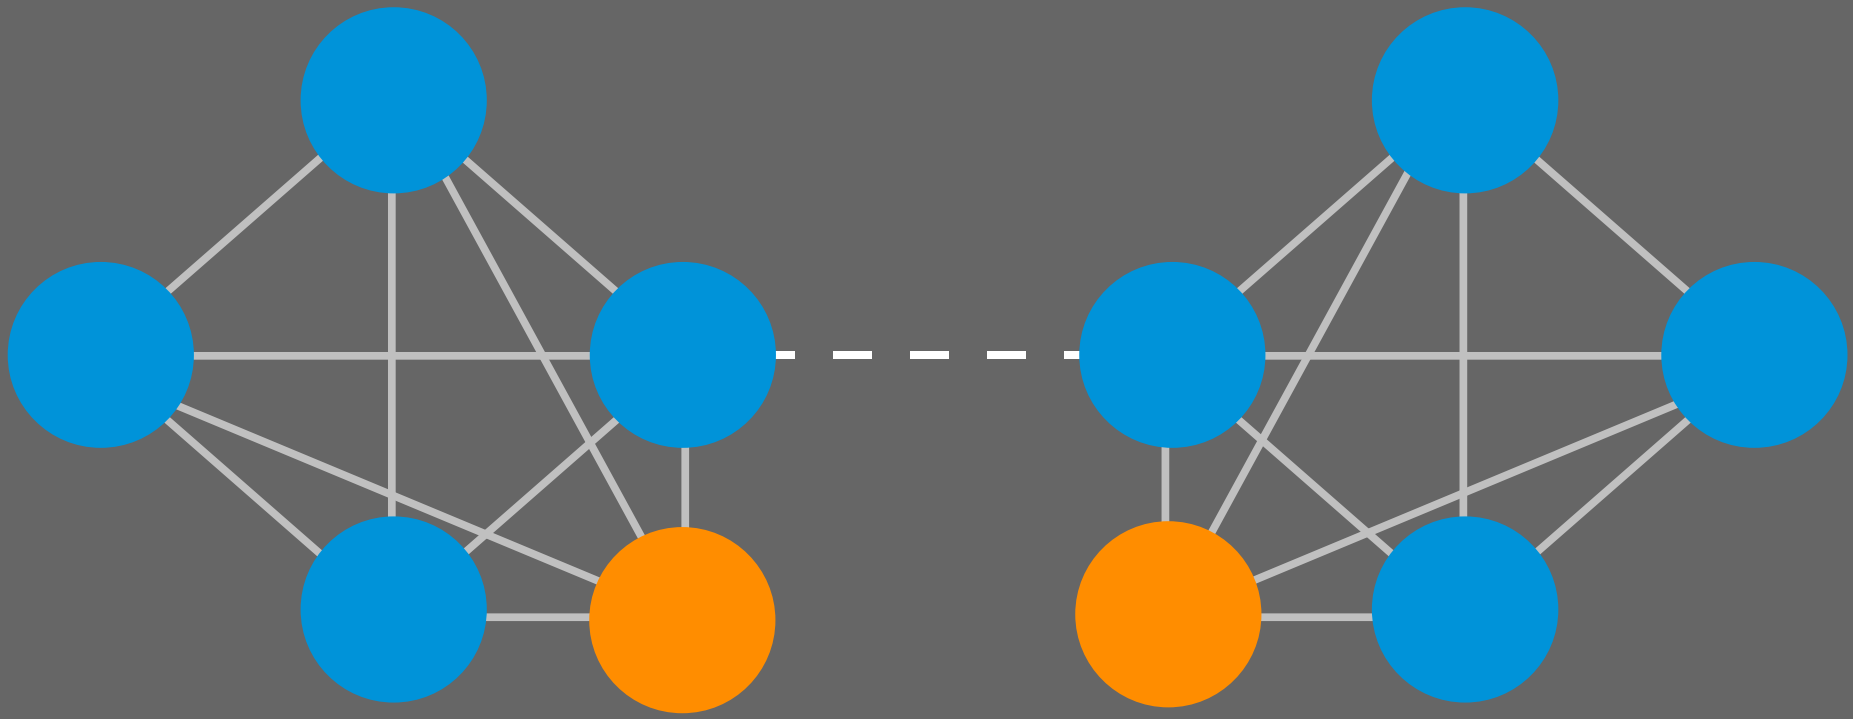

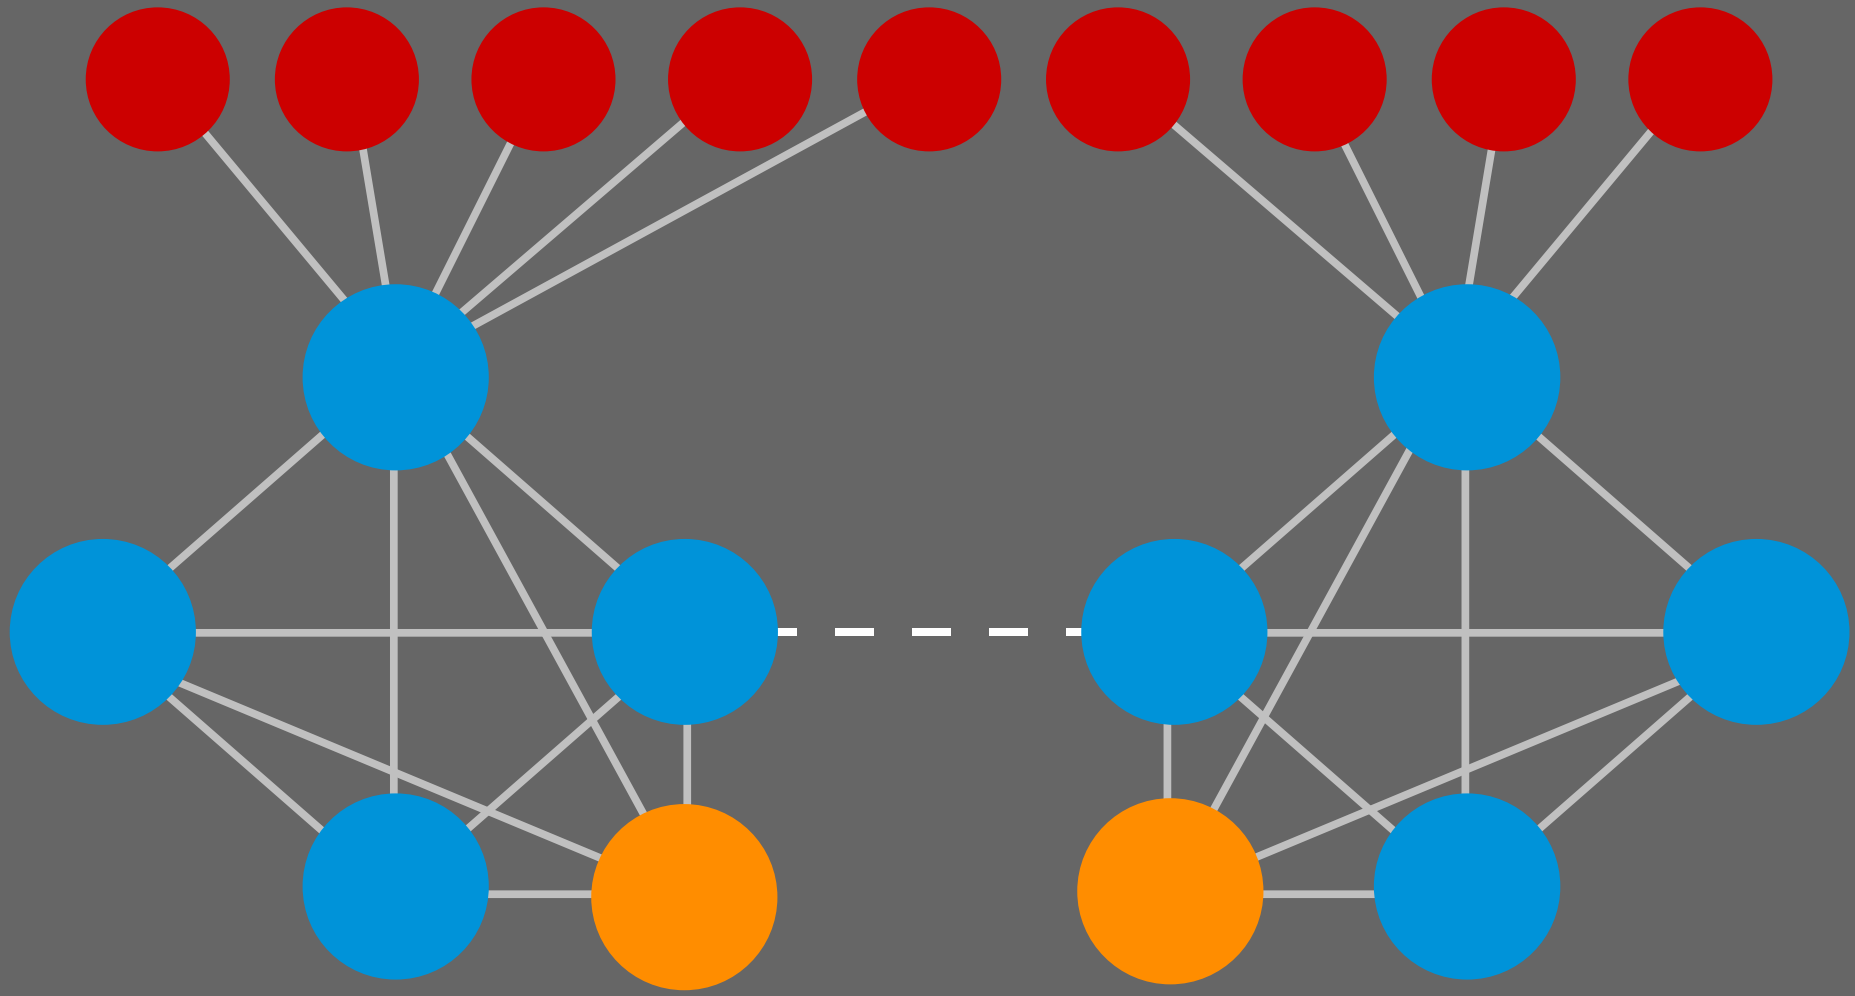

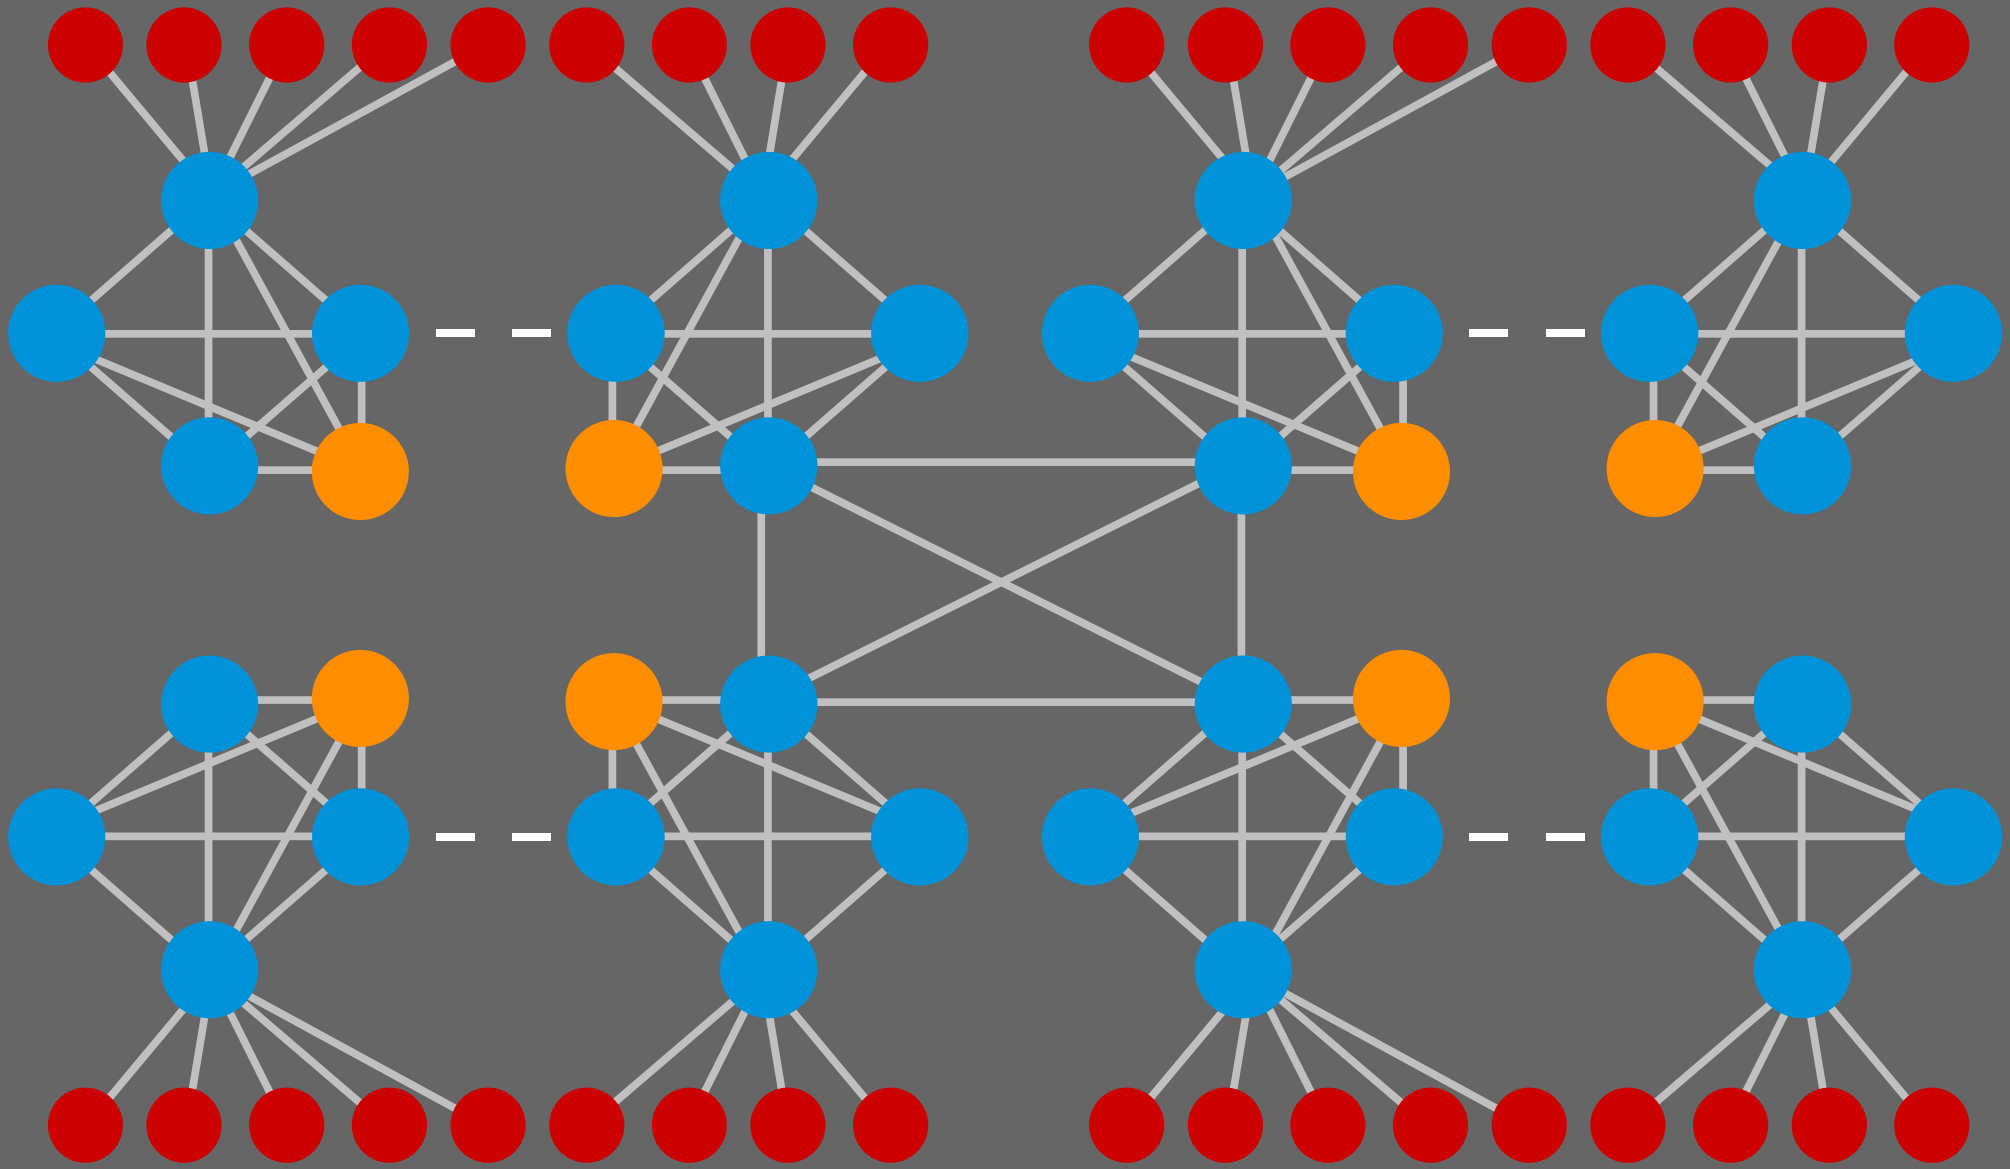

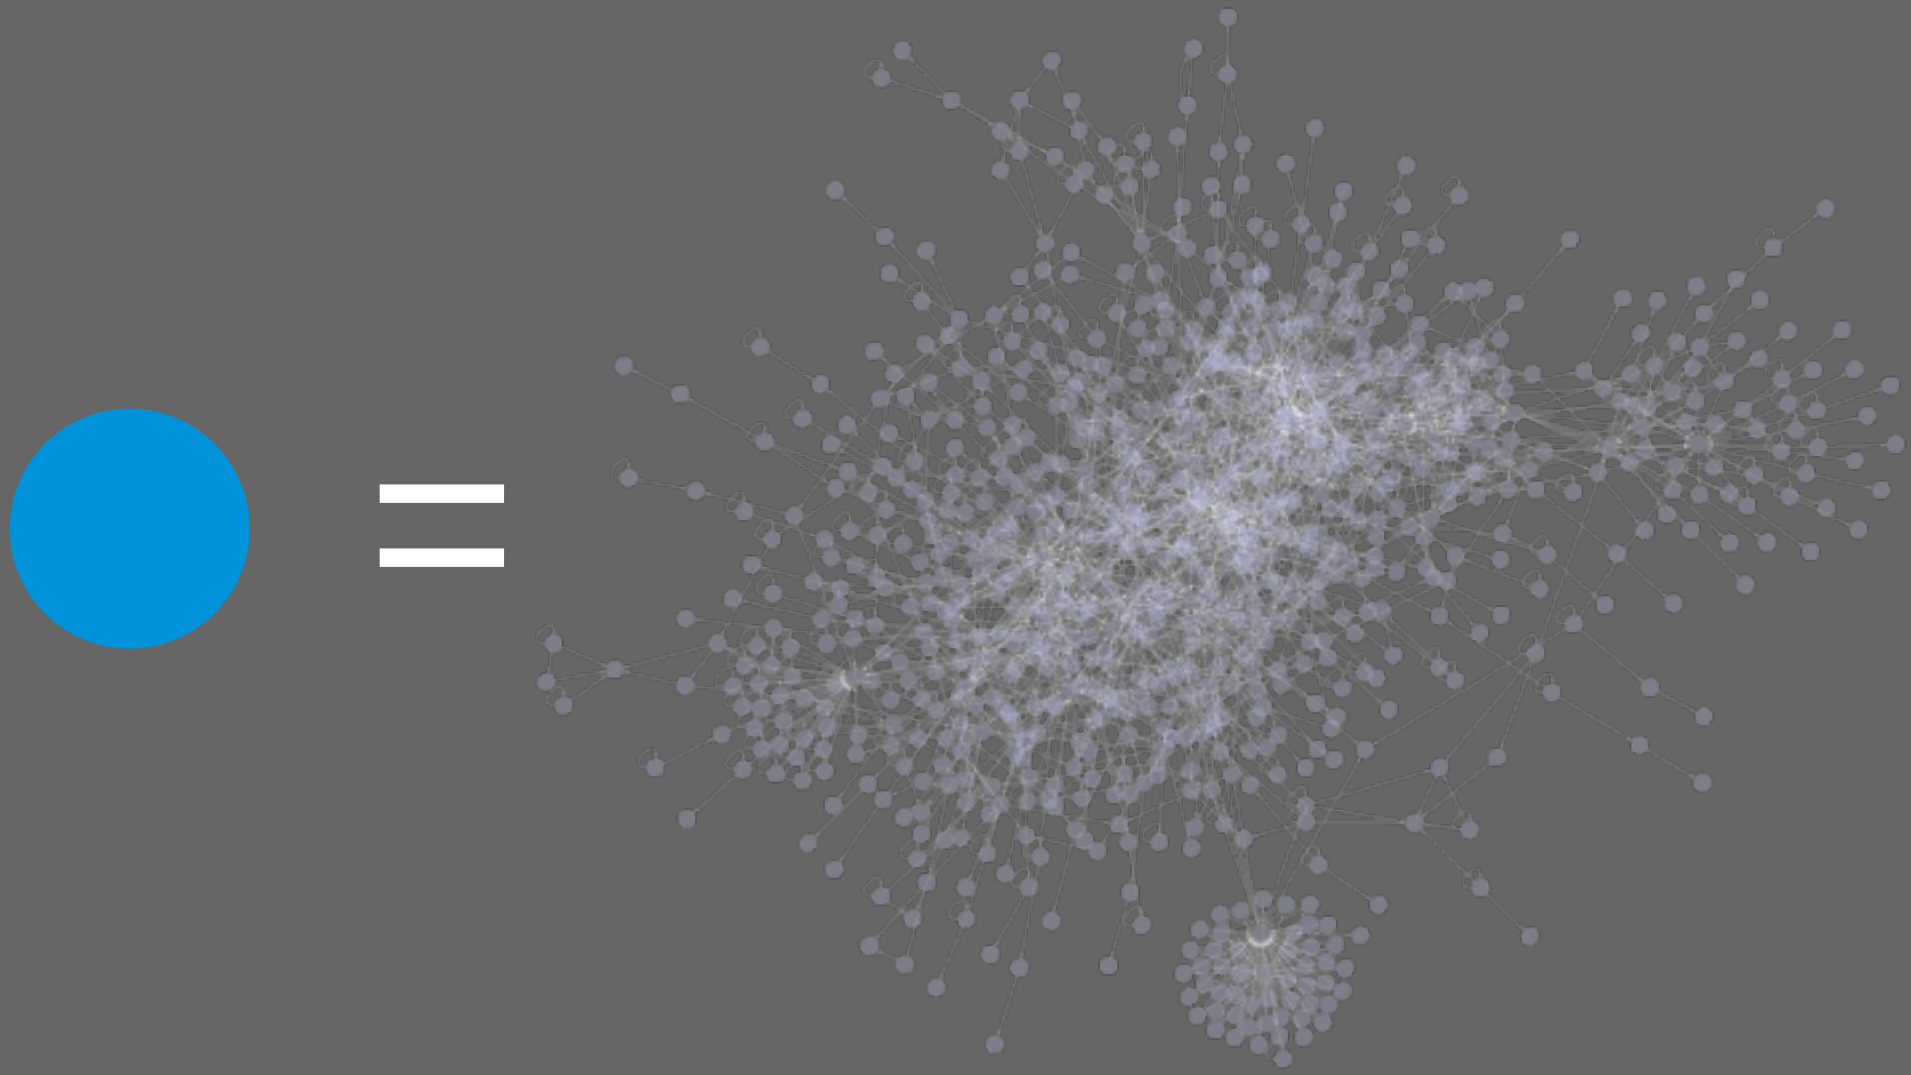

Robustness vs Flexibility

Can't Touch vs Can't Support

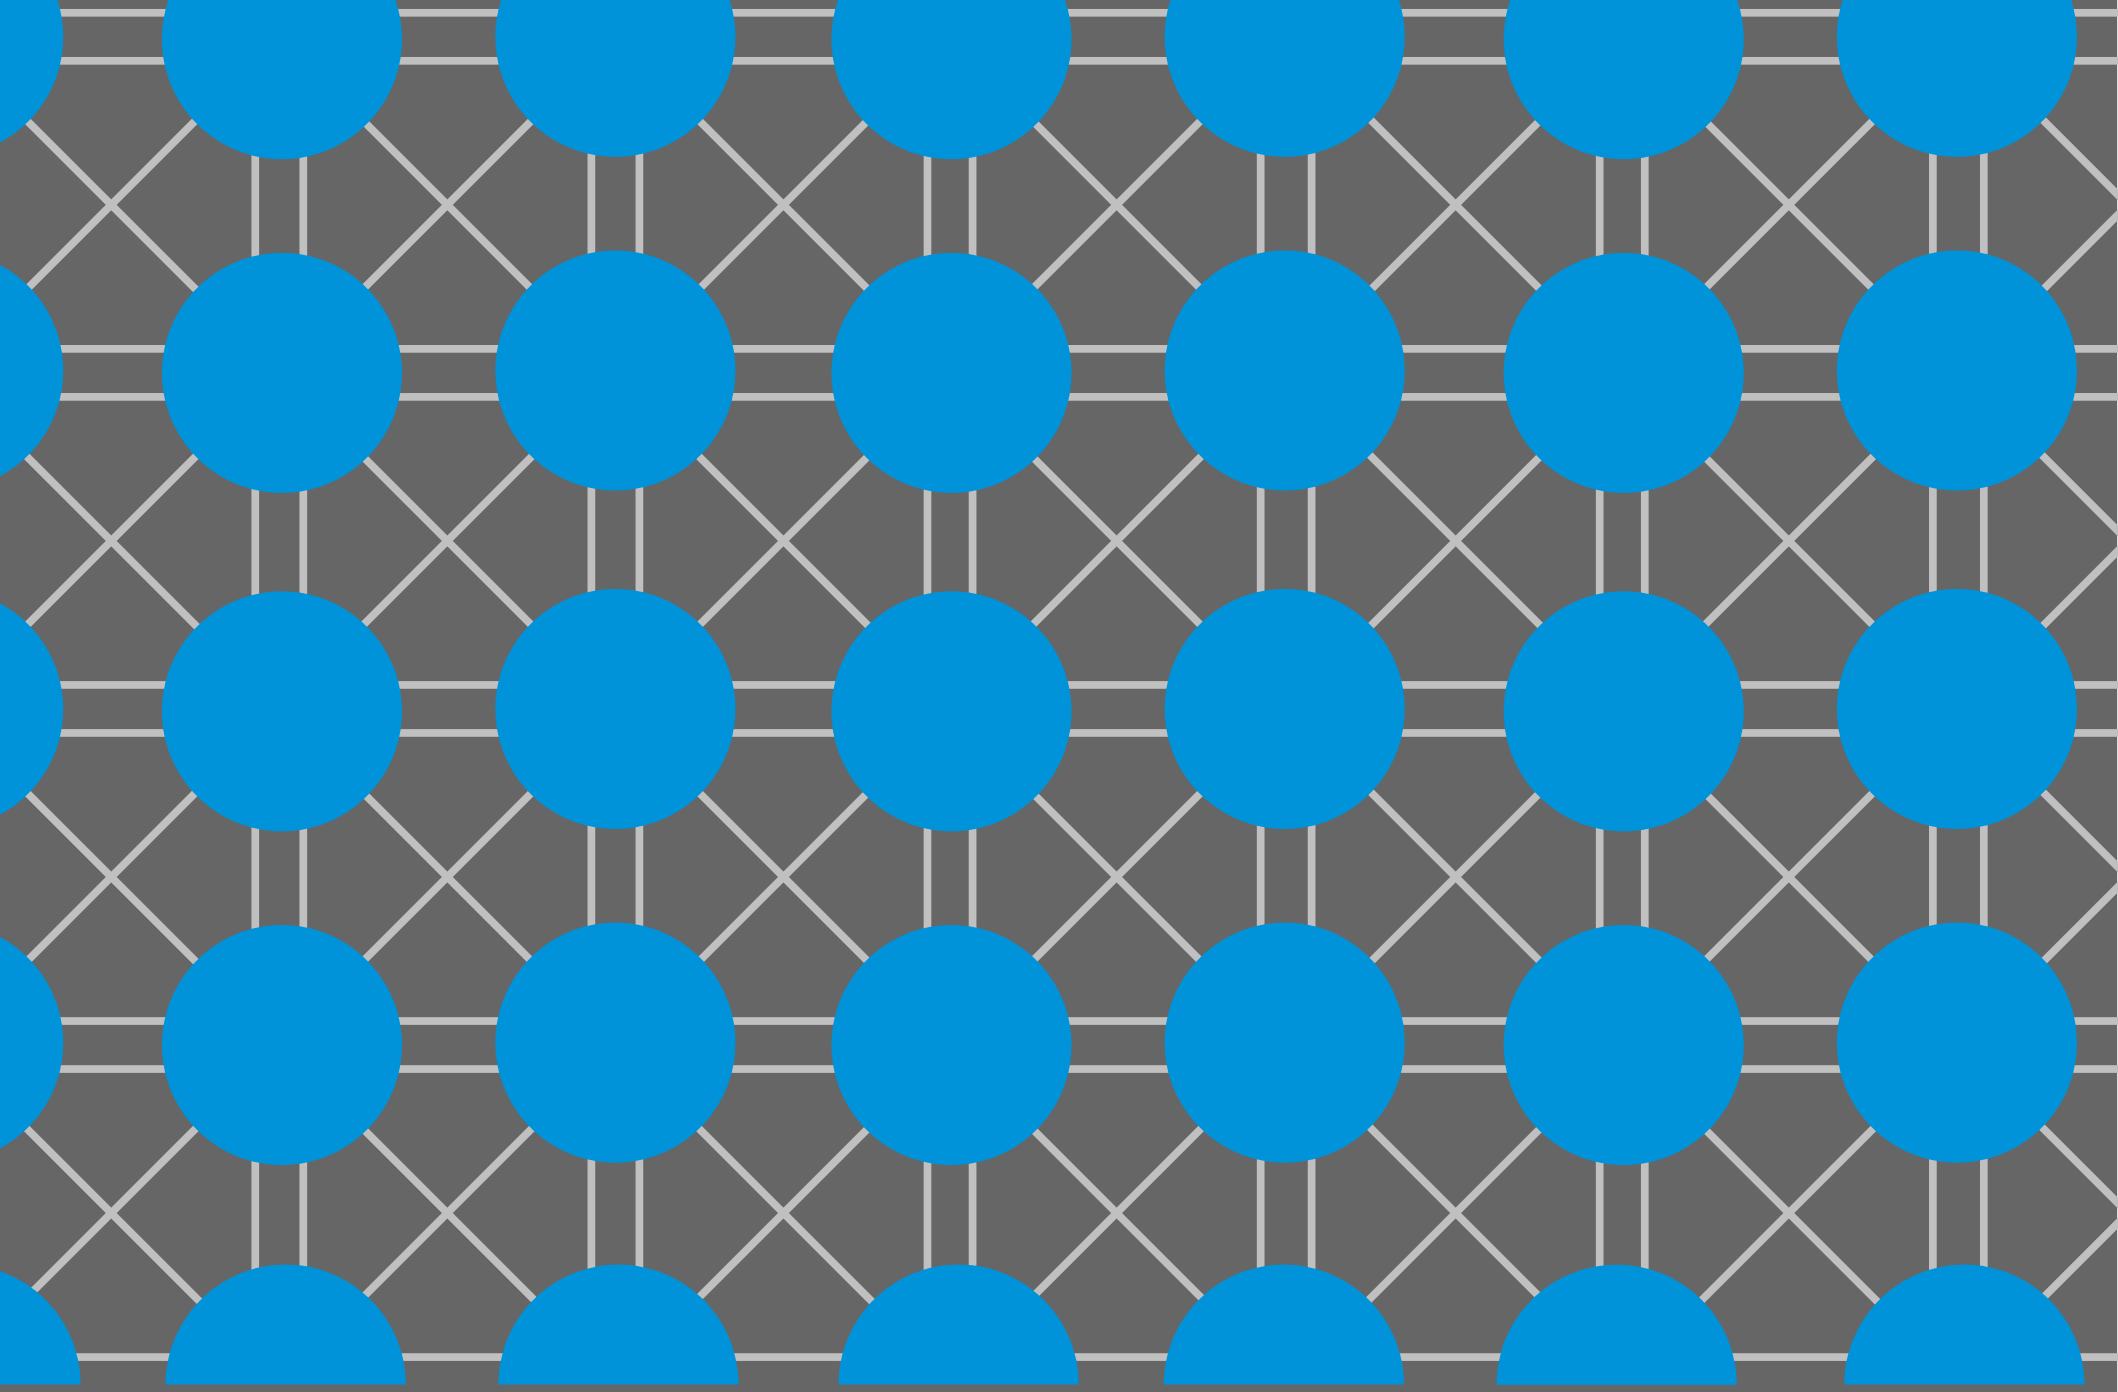

Grid Workload. Any Color You Want

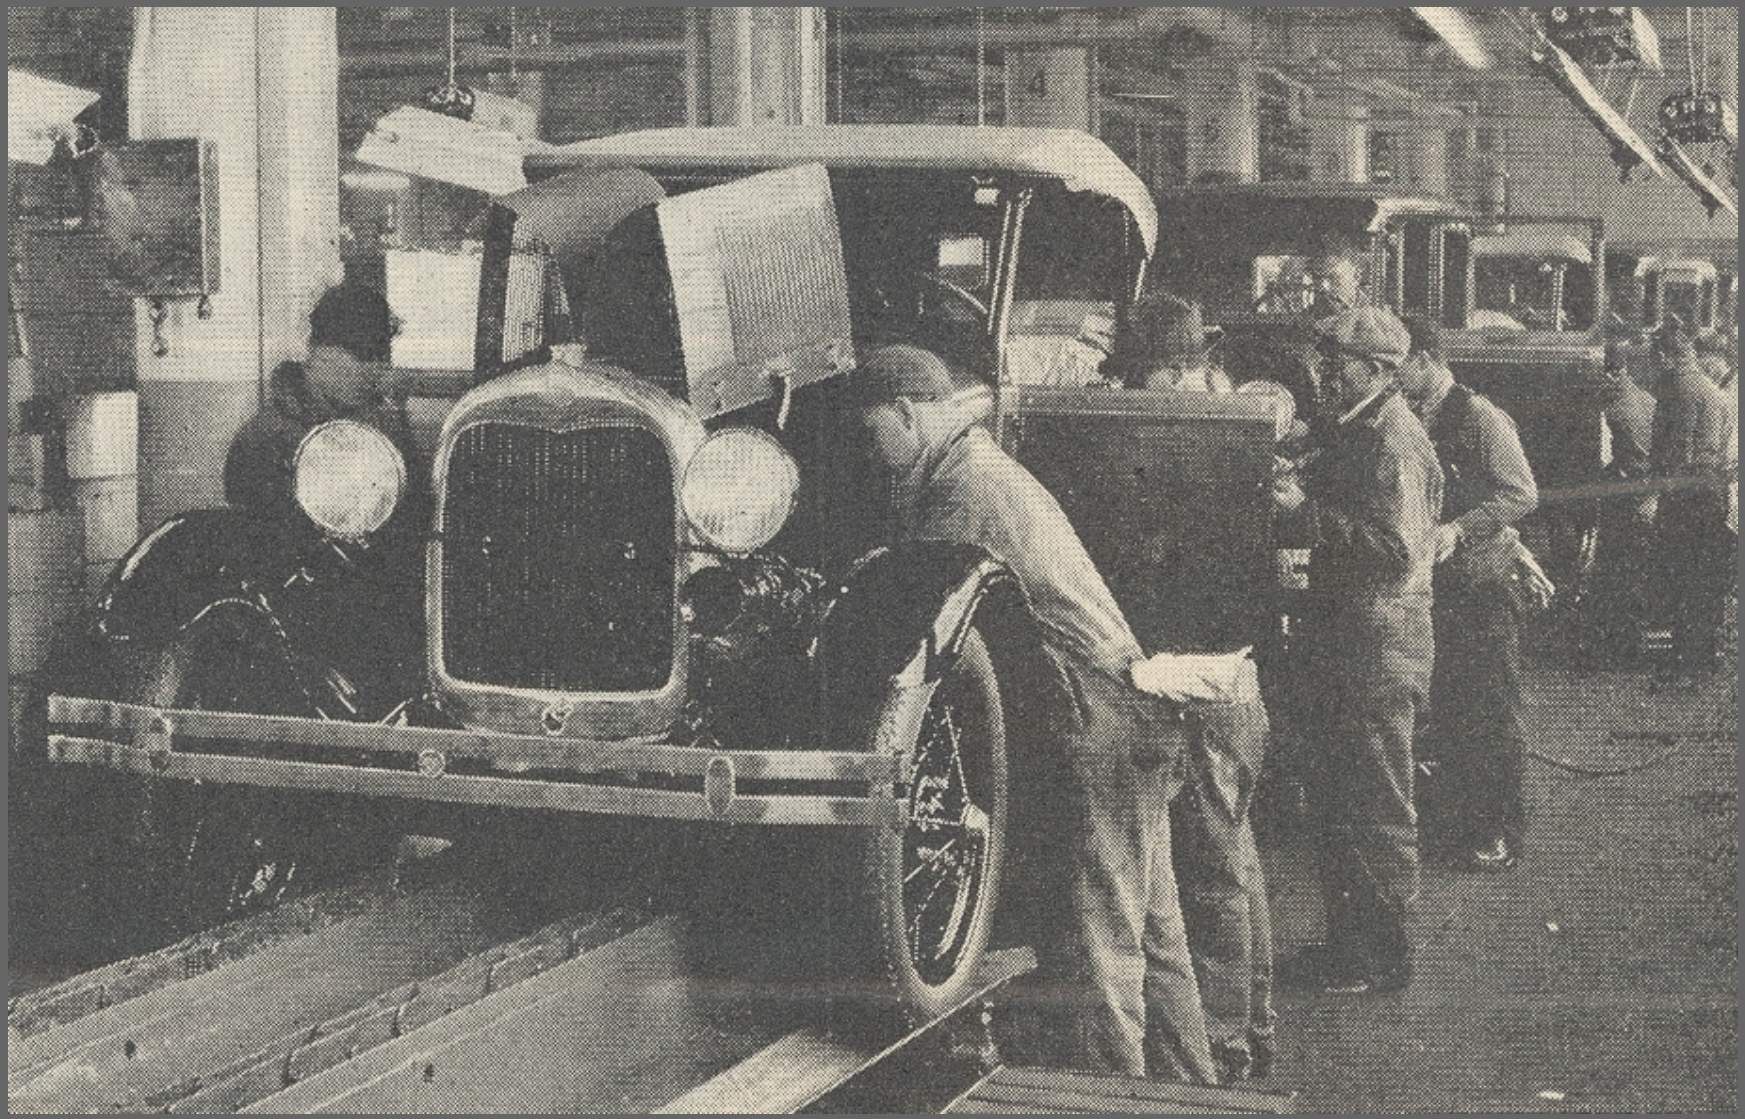

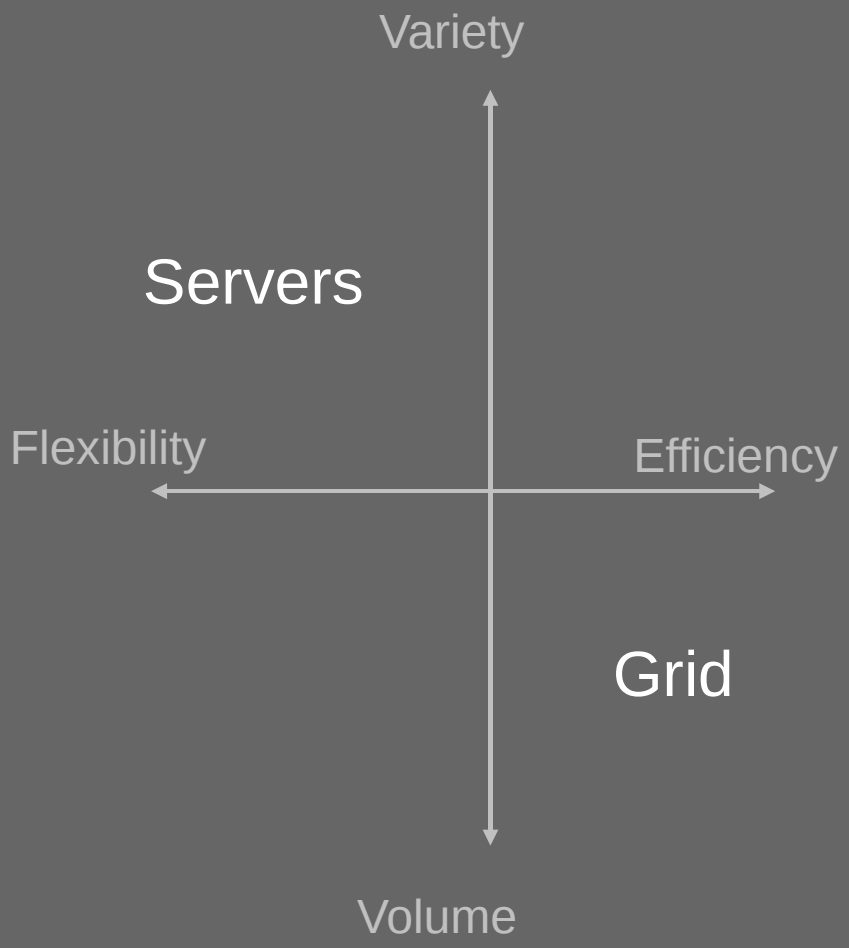

Industry Leaders vs Avg in Business Convergence

12% revenue growth vs. 4%

36% EPS growth vs. 7%

6% higher EBITD margins

4% higher return on equity

8% average higher return on assets

14% higher ROI

http://www.btminstitute.org/

Flexibility > Efficiency

Builds are Workloads

Builds are Workloads

Infrastructure adapts to demand

Many environments simultaneously

Get more workloads on the efficient grid

Requirements

Available Inventory

Rapid re-provisioning

Multiple Standard Builds

MRG Grid is based on the Condor Project created and hosted by the University of Wisconsin, Madison

Condor has a >20-year history and runs many of the largest grids in the world

Red Hat and the University of Wisconsin have signed a strategic partnership around Condor:

- University of Wisconsin makes Condor source code available under OSI-approved open source license
- Red Hat & University of Wisconsin jointly fund and staff Condor development on-campus at the University of Wisconsin

Red Hat and the University of Wisconsin's partnership will:

- Add enhanced enterprise features, management, and supportability to Condor and MRG Grid
- Add High Throughput Computing capabilities to Linux

Red Hat Additions

Enterprise Supportability

- Break out Condor from statically-linked blob to multiple well-maintained and individually patchable rpm's

Web-Based Management Console

- Unified management across all of MRG for job, system, license management, and workload management/monitoring

Low Latency Scheduling

- Enable job submission to Condor via AMQP Messaging clients
- Enable sub-second, low-latency scheduling for sub-second jobs
- Back MS Excel calculations with a grid via MRG C# client

Virtualization Support via libvirt Integration

Cloud Integration with Amazon Ec2

- Enable automatic cloud provisioning, job submission, results storage, teardown via Condor scheduler
- Runs as a job, so it can be a dependency for other jobs or executed based on rules (e.g. add capacity in in the cloud if local grid out of capacity)

Concurrency Limits

- Set limits on how much of a certain resource (e.g. software licenses, db connections) can be used at once

Dynamic Slots

- Mark slots as partitionable and sub-divide them dynamically so that more than one job can occupy a slot at once

MRG Grid Architecture Components

Central Manager: Schedules Jobs

- collector: collects info about pool status
- negotiator: responsible for match-making. Informs submit nodes about execute nodes & vice-versa

Submit Node: Submit Jobs

- schedd: schedules jobs and stores in job queue
- shadow: spawned to manage jobs

Execute Node: Executes Jobs

- startd: enforces policies, spawns job to starter
- starter: process that spawns remote job and sends statistics to submitter

Master Daemon manages other daemons

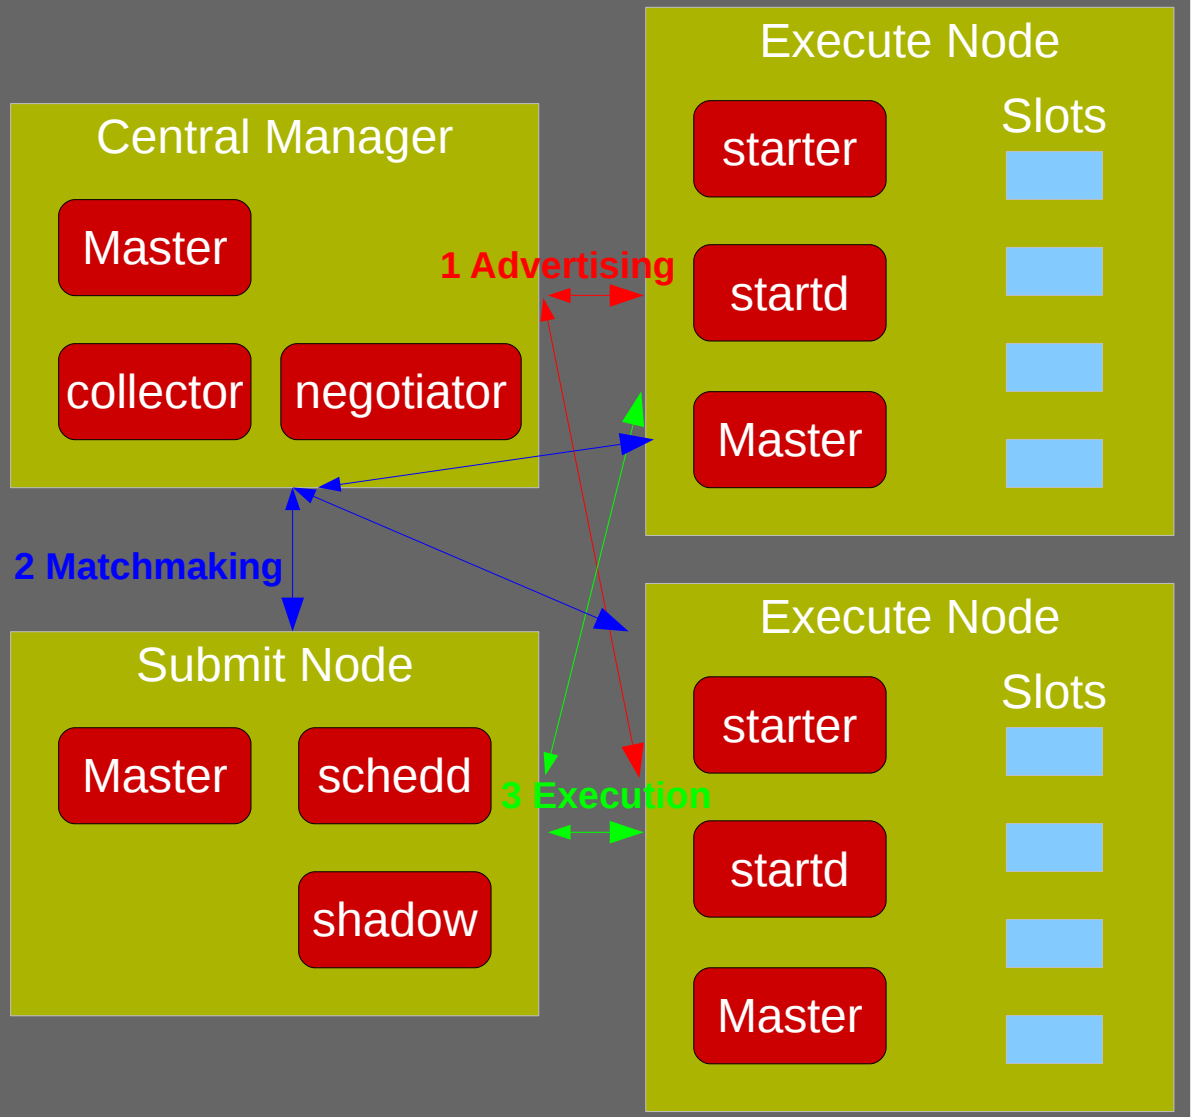

Three Protocols

Advertising – This protocol helps Job and Machine components to find each other.

Matchmaking protocol – used by Central Manager for creating connections between Submit Nodes and Execute Nodes.

Execution Protocol – between Submit Node and Execute Node. Bi-directional connections are made and negotiation of job candidates and file transfers, including executables, performed.

Red Hat MRG – based on Condor

No application changes required

Multiple grid backends

Dependency resolution

Red Hat Network Satellite

Native systems management

Provisioning

Software dependency resolution

Available Inventory

```
[root@emorisse2 ~]# condor_status

Name               OpSys   Arch   State     Activity LoadAv Mem   ActvtyTime

fedorabox          LINUX   INTEL  Unclaimed Idle     2.260   501  0+00:00:07
slot1@emorisse2    LINUX   X86_64 Unclaimed Idle     0.470  1236  0+00:00:04
slot2@emorisse2    LINUX   X86_64 Unclaimed Idle     0.000  1236  0+00:00:05

                     Total Owner Claimed Unclaimed Matched Preempting Backfill

      INTEL/LINUX        1     0       0         1       0          0        0
    X86_64/LINUX         2     0       0         2       0          0        0

            Total        3     0       0         3       0          0        0
```

Base Install

Just enough to get machine up and running

Just enough to attach to the Grid

Custom ClassAd name/value pair:

Provisioning = Available

Build As Workload

```
Requirements   = Arch =="x86_64"  && Provisioning ==
   “Available”
Executable      = /var/lib/condor/job-scripts/mechanize.pl
Log             = /tmp/condor.mechanize.log
output          = /tmp/condor.mechanize.html
should_transfer_files = YES
when_to_transfer_output = ON_EXIT
Queue           100
```

Job Details

Run on the system to be reprovisioned

Logs on to the Satellite Server

Requests live re-provisioning

v0.1: Perl script that interacts with the Satellite Web UI – Satellite version specific

v0.2: Uses new API calls in Satellite 5.3

Architecture Best Practices

- Custom ClassAd
- Create job to “clean up” after usage to restore to base
- Part of usage model
- Or set ClassAd and use cronjob

Flexibility + Efficiency

Get more on the grid

No required code changes

Change the grid base on-demand

Lower time to delivery of service

More Info...

Red Hat MRG – www.redhat.com/mrg

MRG docs – www.redhat.com/docs/en-US/Red_Hat_Enterprise_MRG/

Condor Project – www.cs.wisc.edu/condor

Red Hat Network Satellite – www.redhat.com/rhn

Spacewalk – www.redhat.com/spacewalk

QUESTIONS?

TELL US WHAT YOU THINK:
REDHAT.COM/SUMMIT-SURVEY