



DREAMWORKS
— ANIMATION SKG —

DreamWorks Animation SKG

Only studio to release 3 CG animated features
in one year

Consistently in top 5 each year

Collaboration and Innovation with partners

Scale of Digital Resources

Each Film Takes 4-5 Years To Make

200+ Workstations

65+ Million CPU Hours

Peak use of 15,000 cores of 22,000 total cores

200+ TB's of data

500+ Million files

10Gb/s WAN, 12% of Mad3 rendered remotely

Imagine Having a **DOZEN Films in Production** at ONCE!

Renderfarm Scale

Hundreds of thousands of jobs per night

Up to 20,000 jobs running concurrently

Deadline-based turnaround on priority order

Average runtime ~ 15 min, >90% utilization

DreamWorks And Open Source

DWA helped **lead industry** toward Linux and open source

Enterprise class solution - Linux-based HP workstations

With Red Hat we helped to enable **commodity computing** for our industry...which still uses proprietary grid middleware

Expensive in \$\$\$ and effort

We followed the **open source path** for our grid middleware

Renderfarm Requirements

Dynamic Slots	HGQ	Round-Robin Rate
Job Group Priority Order	FCFS (PostJobPrio1)	Progressive Rendering
Job Policies and Steering	Resource Limits	Suspend/ Continue
Schedd Scaleout	Multi-target VIEW_HOST	Negotiator & Schedd Stats

Looking Forward

Improved schedd HA	Multiple default concurrency limits
Startd-local limits	Cgroups & FS namespace isolation
"Claim Partitionable Leftovers" - but....	Plumage & add'l statistics


Operational Visibility


Our artists need much more than a "black box"


Built our own management **tailored for rendering** workflow


Built a **realtime cache** of farm data (stock + our metadata)
inclusive of unreleased / dependent jobs


Using provided statistics, finding areas of **improvement**

















Lessons Learned From Neg Stats

Consistent Matches = Duration-bound

Need deeper stats around Rejections, why 2X?

No reschedules/drop interval - but verify knobs

Phase1Duration has become 40%+ of Duration

Workflow Improvements

DAGman Service Interface

- Query submitted AND dependent jobs

- Live editing (priority, kill a graph node)

Accounting group role-based access

- Team-lead manual (re)prioritization

- Deadline-oriented workflow enablement

DAG meta dependencies

Thank You!

Questions?

Lans Carstensen

Lans.Carstensen@dreamworks.com

Principal Engineer, Digital Operations