

SUMS

**STAR Unified Meta Scheduler
Over view, improvements
and research**

Levente B. Hajdu

Overview

- Why are STAR members encouraged to use SUMS ?
- Improvements and additions to SUMS
- Research
 - Job scheduling with load monitoring tools
 - Request Definition Language (RDL)
- Who contributes to SUMS research and development ?



Intro

- SUMS is a scheduler written mainly by STAR with contributions from many other organizations. It's used to run programs on datasets (large set of files) that may be distributed across different nodes, clusters, batch systems, and, sites. It is distributed free of charge under the GNU General Public License.
- SUMS unifies scientific computing requests into one general format.
- It is hoped that more experiments will join us and use our scheduler.



Why are STAR members encouraged to use SUMS ?

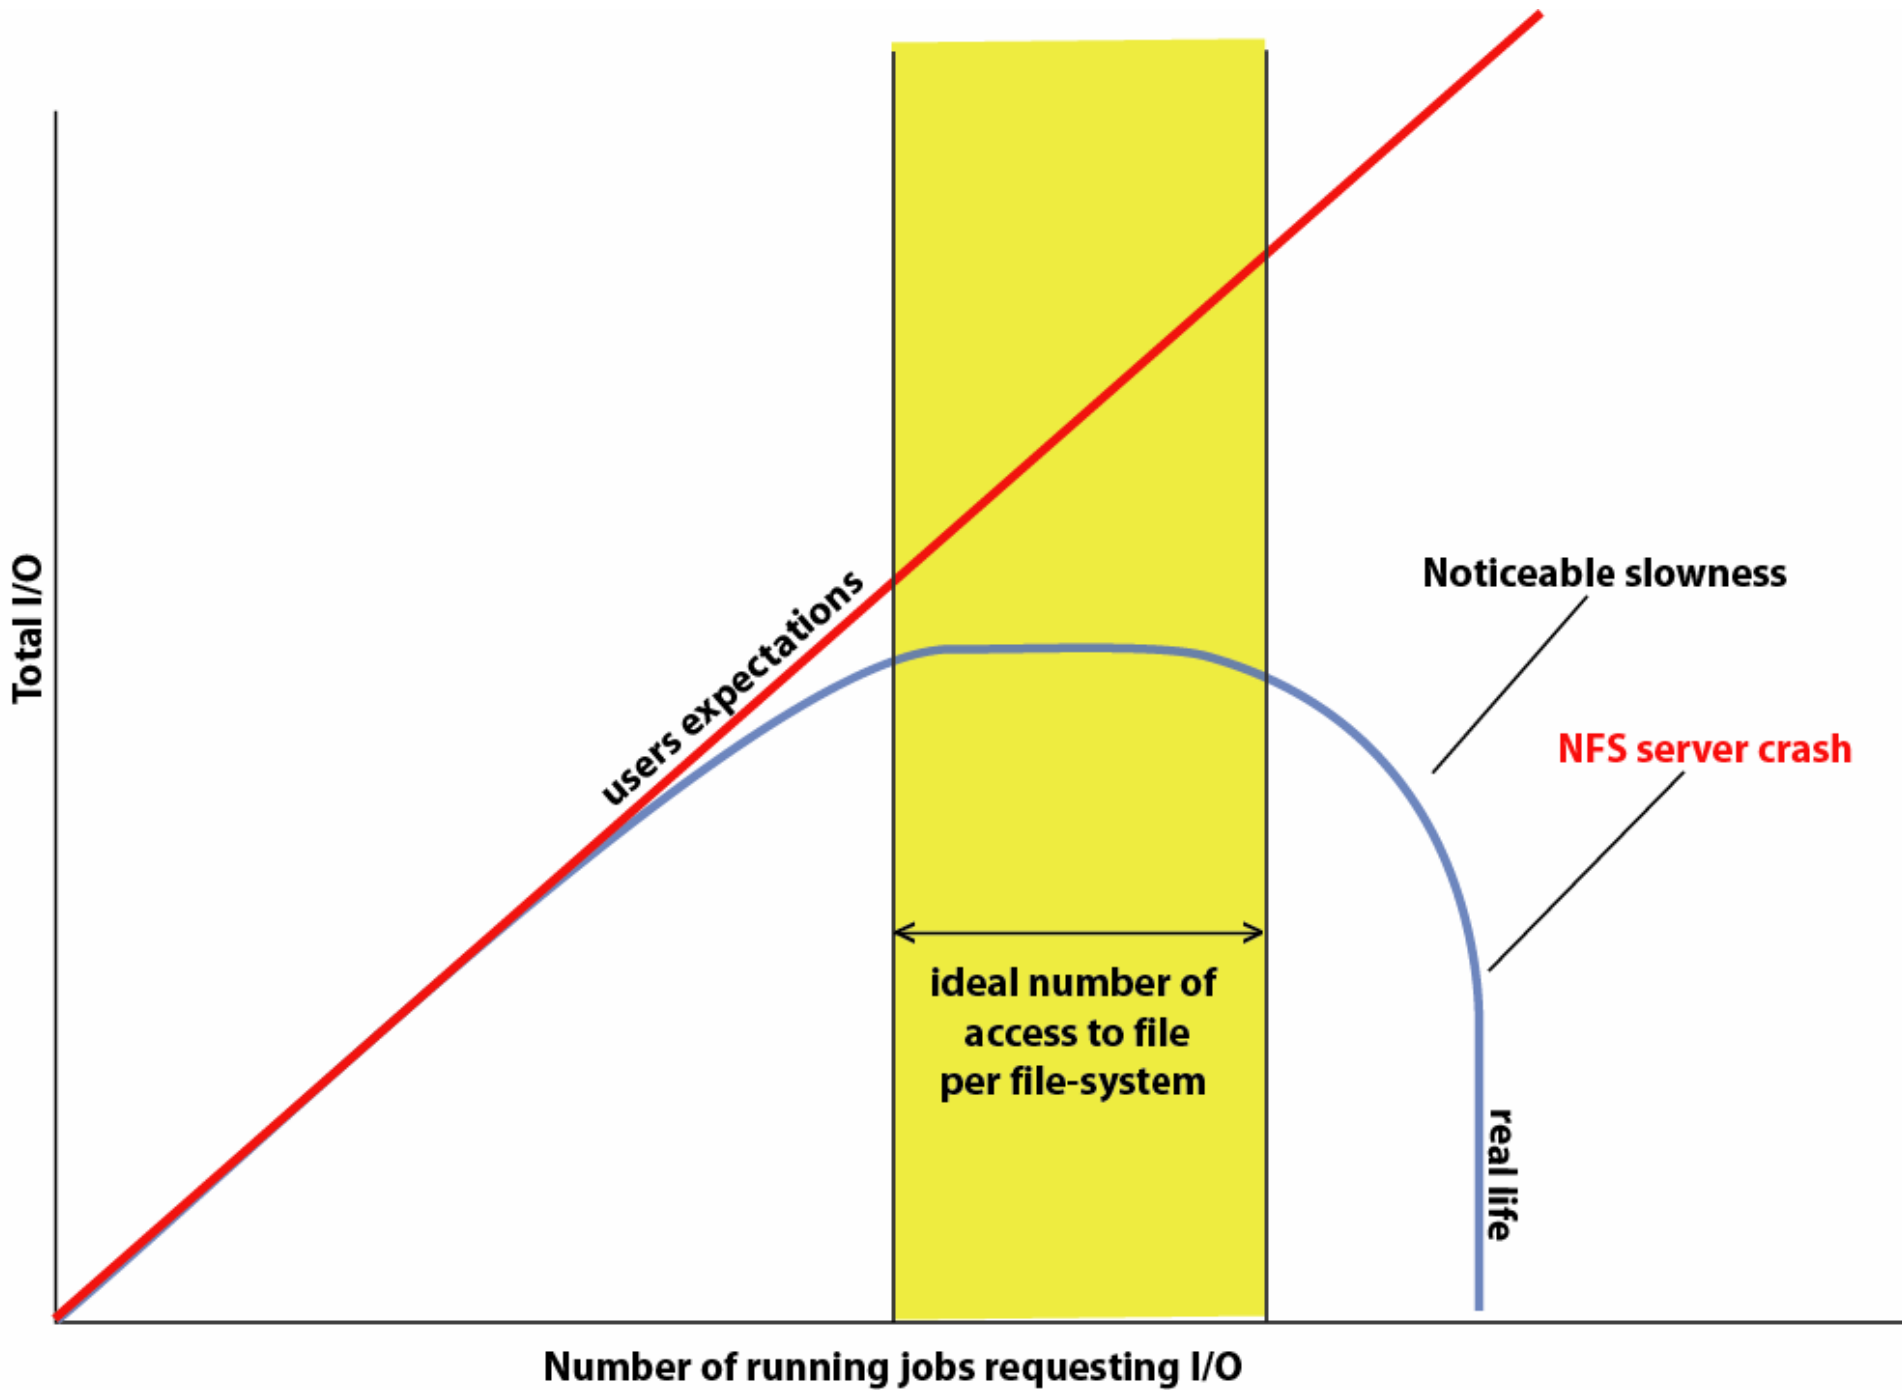
- Resolves Problem of over use of limited resources, to the point of failure.
- Resolves Problem of accessing files on distributed disks
- Breaks dependents on any one particular batch system
- Is helpful in file catalog resolution and optimizing jobs for less experienced users




Why are STAR members encouraged to use SUMS ?

- Resolves Problem of over use of limited resources, to the point of failure.
- Resolves Problem of accessing files on distributed disks
- Breaks dependents on any one particular batch system
- Is helpful in file catalog resolution and optimizing jobs for less experienced users

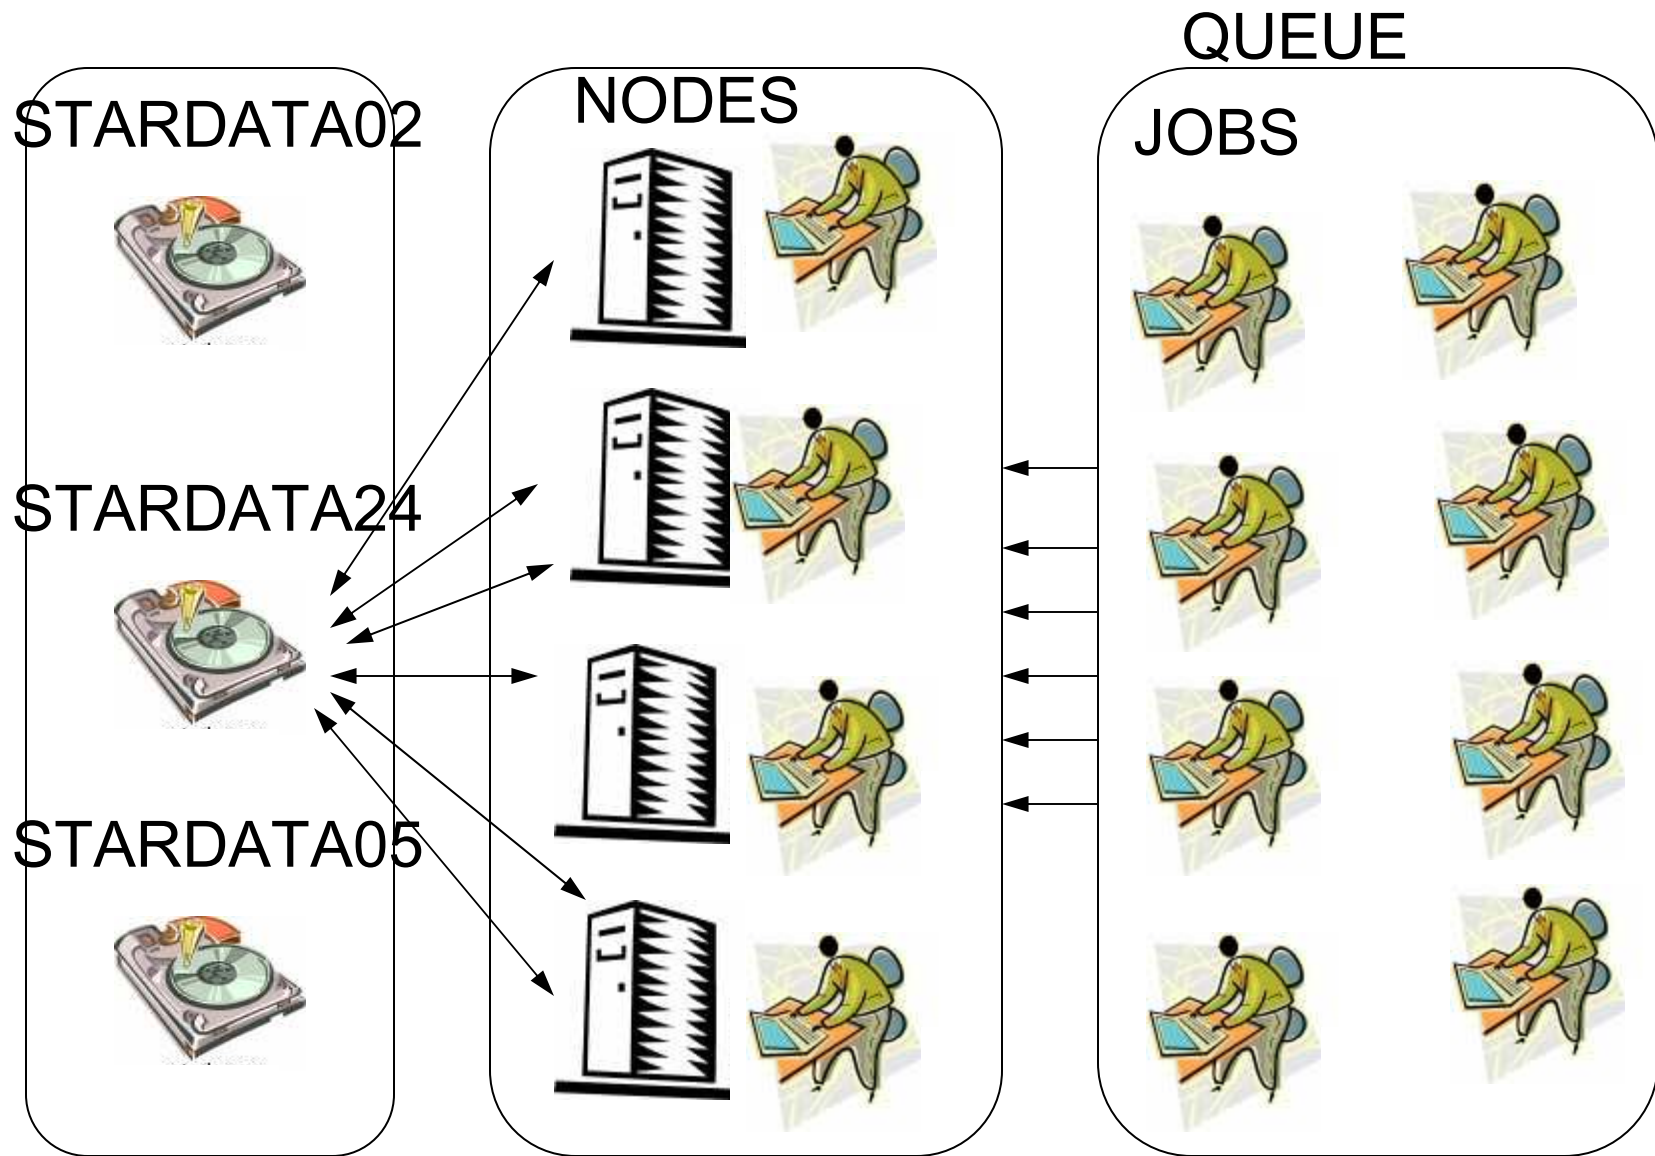




RCF



RCF



Queued jobs



SD05 = 800
SD24 = 50



SD05 = 750



SD05 = 700
SD02 = 450

Running jobs



SD05 = 500
SD24 = 50



SD05 = 500
SD24 = 10



SD05 = 1000
SD24 = 2



SD05 = 30
SD24 = 3



SD05 = 30
SD02 = 600

SD05 virtual resource
(2040 units total)

SD24 virtual resource
(102 units total)

SD02 virtual resource
(800 units total)

STARDATA05

STARDATA24

STARDATA02



- **Past reactions demonstrate that without proper explanation users believe this kind of strategy slows down their jobs. In reality the strategy (well implemented) greatly increases farm through-put.**
- **Jobs that need access to resources no one else is using run first (and fast).**
- **Jobs don't need to compete for the same resource, slowing each other down, running for longer than they have to (slots open up faster).**
- **My jobs are not run sequentially they are run in the most optimal order.**



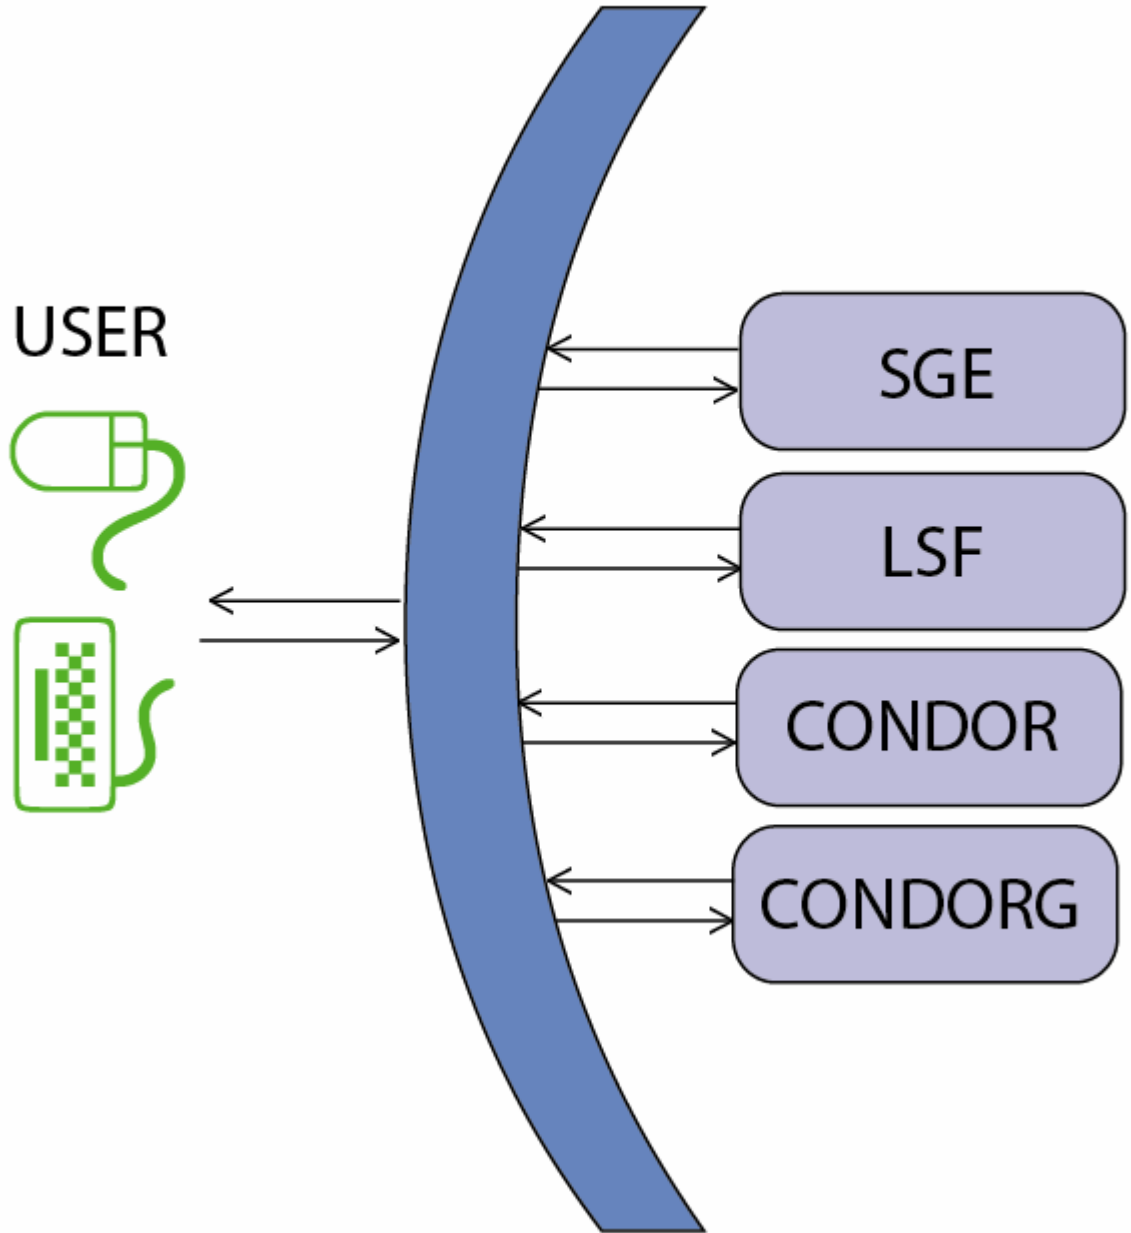
Why are STAR members encouraged to use SUMS ?

- Resolves Problem of over use of limited resources, to the point of failure.
- Resolves Problem of accessing files on distributed disks
- Breaks dependents on any one particular batch system
- Is helpful in file catalog resolution and optimizing jobs for less experienced users

Why are STAR members encouraged to use SUMS ?

- Resolves Problem of over use of limited resources, to the point of failure.
- Resolves Problem of accessing files on distributed disks
- Breaks dependents on any one particular batch system
- Is helpful in file catalog resolution and optimizing jobs for less experienced users

sums



Why are STAR members encouraged to use SUMS ?

- Resolves Problem of over use of limited resources, to the point of failure.
- Resolves Problem of accessing files on distributed disks
- Breaks dependents on any one particular batch system
- Is helpful in file catalog resolution and optimizing jobs for less experienced users

Improvements and additions to SUMS

- **Resubmit\Status\kill switches**
- **minMemory, maxMemory, minStorageSpace, maxStorageSpace keywords**
- **SGE dispatcher**
- **\$FILEBASENAME**
- **schema validation**
- **Extend report file**
- **uuidgen**



Resubmit\status\kill switches

- **star-resubmit -kr all DD9CFB586F4139E8D14C6.session.xml**
//Kill and resubmit all jobs
- **star-resubmit -r 1,2,3,5 DDD9CFB586F4139E8D14C6.session.xml**
//Resubmit jobs 1,2,3, and 5
- **star-resubmit -s 5 DD9CFB586F4139E8D14C6.session.xml**
//Get the status of job 5



maxMemory and maxStorageSpace

- The “max” limits lets sums know not to submit your job in a queue where memory or storage space may be an issue
- Under LSF it also lets the queue know to kill the job if it exceeds the max limits.
 - (very handy if you suspect potential for “run away jobs”)

minMemory, minStorageSpace,

- The “min” limits let the batch system know how much space should be freed before you job can run
- Under LSF, if this limit is smaller then the default limit your job will most likely find a slot before other jobs seeking larger slots.
- Under LSF, if this limit is way to high your job may never find a slot (may never run)



The SGE dispatcher

- PDSF will be converting from the LSF batch system to the SGE batch system. At the request of PDSF a new dispatcher plug-in was written for the SGE batch system.
- During the transition period, as LSF nodes go down and SGE nodes come up, SUMS will determine which batch system local files are on and format submits in the appropriate way.
- Our new SGE dispatcher plug-in lets SUM dispatch to SGE just like it did previously with LSF.
- Users at PDSF will notice **absolutely no** difference in how they submit there jobs with SUMS.

Schema validation \ Extend report file

- Users request are compared to a JDL schema. If there is an error a validation package will give an error stating both the line number and a description of the error. Similar to what you would expect from the java compiler.



TAR scheduler 1.7.X
 File generated at : Fri Feb 11 12:14:48 EST 2005
 This file is a detailed report of SUMS meta job 9990F27153B2088952DFD6E2604917A3 and its 8 processes.

Jobs table :

ID	Queue	Files	estRunTime	Target	submitTime	Dispatcher	successful	Dt
9990F27153B2088952DFD6E2604917A3_0	star_cas_big	4	240.0min	rcas6072.rcf.bnl.gov	Fri Feb 11 12:14:47 EST 2005	LSFDispatcher	YES	40
9990F27153B2088952DFD6E2604917A3_1	star_cas_big	4	240.0min	rcas6072.rcf.bnl.gov	Fri Feb 11 12:14:43 EST 2005	LSFDispatcher	YES	40
9990F27153B2088952DFD6E2604917A3_2	star_cas_big	4	240.0min	rcas6072.rcf.bnl.gov	Fri Feb 11 12:14:39 EST 2005	LSFDispatcher	YES	40
9990F27153B2088952DFD6E2604917A3_3	star_cas_big	4	240.0min	rcas6072.rcf.bnl.gov	Fri Feb 11 12:14:34 EST 2005	LSFDispatcher	YES	50
9990F27153B2088952DFD6E2604917A3_4	star_cas_big	4	240.0min	rcas6072.rcf.bnl.gov	Fri Feb 11 12:14:29 EST 2005	LSFDispatcher	YES	30
9990F27153B2088952DFD6E2604917A3_5	star_cas_big	4	240.0min	rcas6072.rcf.bnl.gov	Fri Feb 11 12:14:26 EST 2005	LSFDispatcher	YES	40
9990F27153B2088952DFD6E2604917A3_6	star_cas_big	3	180.0min	rcas6072.rcf.bnl.gov	Fri Feb 11 12:14:22 EST 2005	LSFDispatcher	YES	40
9990F27153B2088952DFD6E2604917A3_7	star_cas_big	3	180.0min	rcas6072.rcf.bnl.gov	Fri Feb 11 12:14:17 EST 2005	LSFDispatcher	YES	30

Request object report :

ID	Name	Type	TimeLimit	MaxMem	Local	S.O.P.	cluster
localQueue	star_cas_dd	LSF	90min	440MB	Y	1	rcas.rcf.bnl.gov
nfsQueue	star_cas_short	LSF	90min	440MB	N	1	rcas.rcf.bnl.gov
longQueue	star_cas_big	LSF	14400min	none	N	100	rcas.rcf.bnl.gov

 * This is a history of Job assignments. *
 * If your jobs are going to the wrong queue, this may tell you why. *

rcas6072.rcf.bnl.gov: nFiles 30 - nProc 8 - valid

Job object 9990F27153B2088952DFD6E2604917A3_0 report :
 9990F27153B2088952DFD6E2604917A3_0 trying star_cas_dd(localQueue)
 9990F27153B2088952DFD6E2604917A3_0 did not fit in star_cas_dd Queue time limit = 90min job estimated run time : 240.0min = (4files/1.0FilesPerHour*60)
 9990F27153B2088952DFD6E2604917A3_0 trying star_cas_short(nfsQueue)
 9990F27153B2088952DFD6E2604917A3_0 did not fit in star_cas_short Queue time limit = 90min job estimated run time : 240.0min = (4files/1.0FilesPerHour*60)
 9990F27153B2088952DFD6E2604917A3_0 trying star_cas_big(longQueue)
 9990F27153B2088952DFD6E2604917A3_0 assigned to star_cas_big(longQueue)

Job object 9990F27153B2088952DFD6E2604917A3_1 report :
 9990F27153B2088952DFD6E2604917A3_1 trying star_cas_dd(localQueue)
 9990F27153B2088952DFD6E2604917A3_1 did not fit in star_cas_dd Queue time limit = 90min job estimated run time : 240.0min = (4files/1.0FilesPerHour*60)
 9990F27153B2088952DFD6E2604917A3_1 trying star_cas_short(nfsQueue)
 9990F27153B2088952DFD6E2604917A3_1 did not fit in star_cas_short Queue time limit = 90min job estimated run time : 240.0min = (4files/1.0FilesPerHour*60)
 9990F27153B2088952DFD6E2604917A3_1 trying star_cas_big(longQueue)
 9990F27153B2088952DFD6E2604917A3_1 assigned to star_cas_big(longQueue)

Job object 9990F27153B2088952DFD6E2604917A3_2 report :
 9990F27153B2088952DFD6E2604917A3_2 trying star_cas_dd(localQueue)
 9990F27153B2088952DFD6E2604917A3_2 did not fit in star_cas_dd Queue time limit = 90min job estimated run time : 240.0min = (4files/1.0FilesPerHour*60)
 "sched9990F27153B2088952DFD6E2604917A3.report" 116L, 8210C

uuidgen

- Universally Unique Identifier Generator
- A UUID can reasonably be considered unique among all UUIDs created on the local system, and among UUIDs created on other systems in the past and in the future.
- Prevents jobs from colliding in the local environment and on the grid



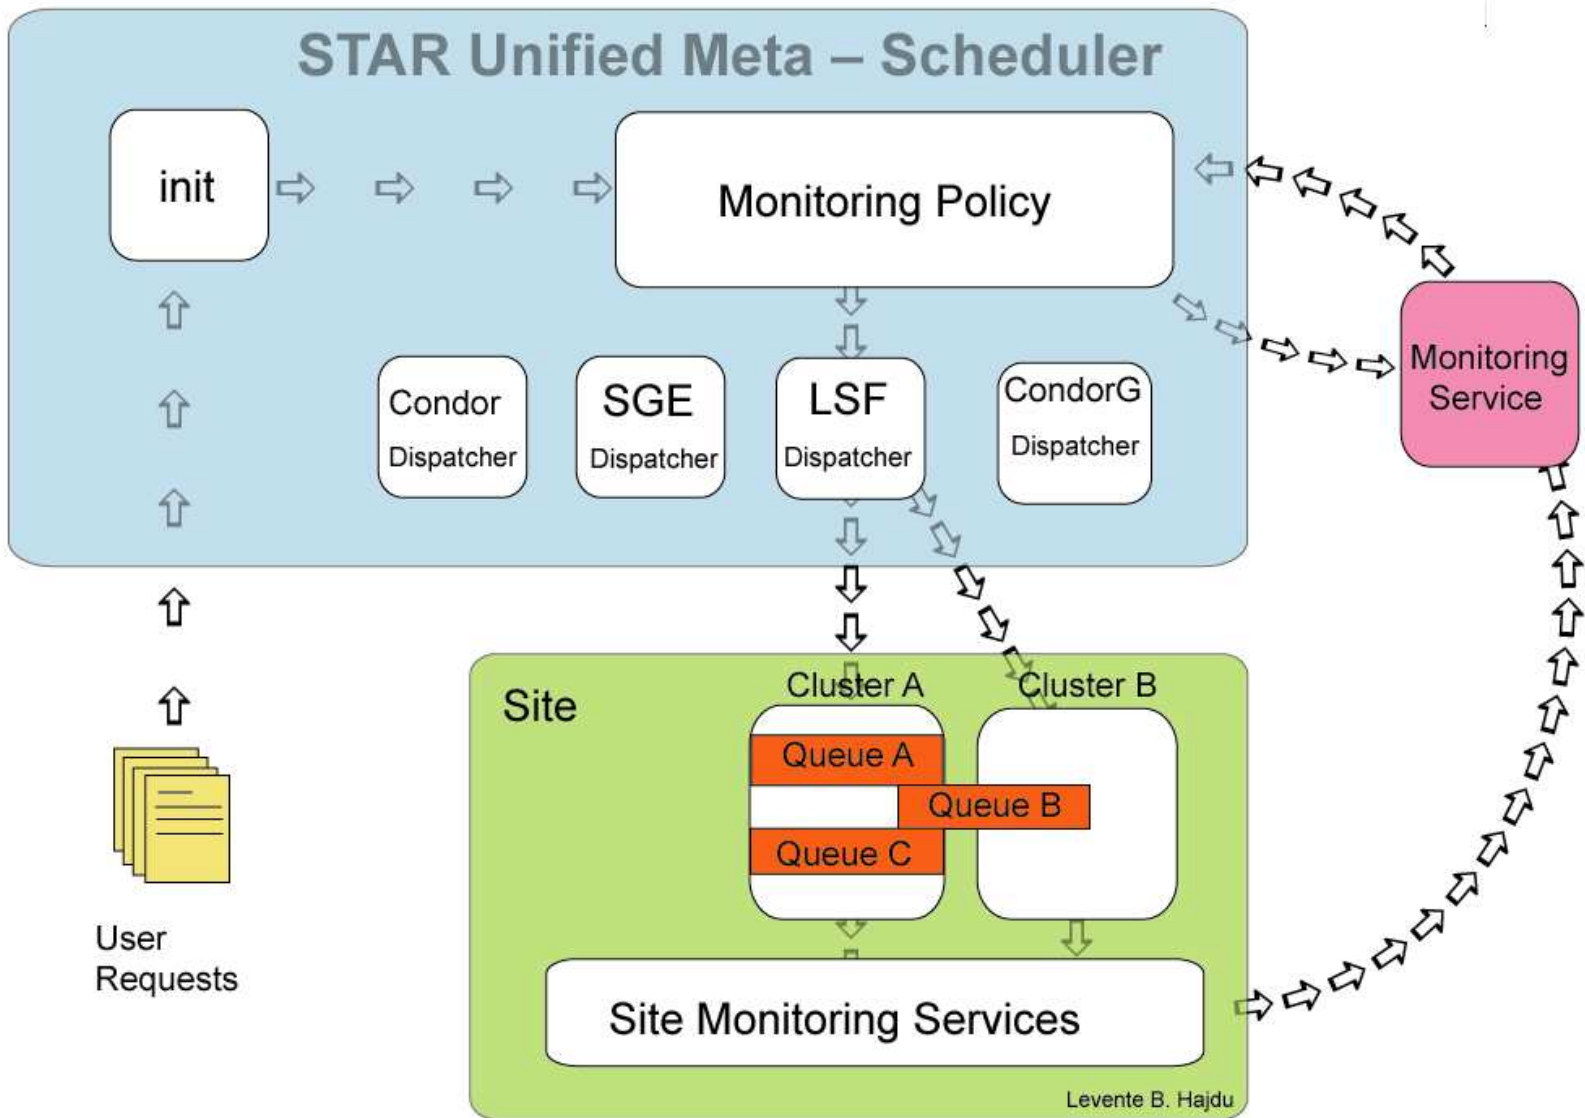


Research

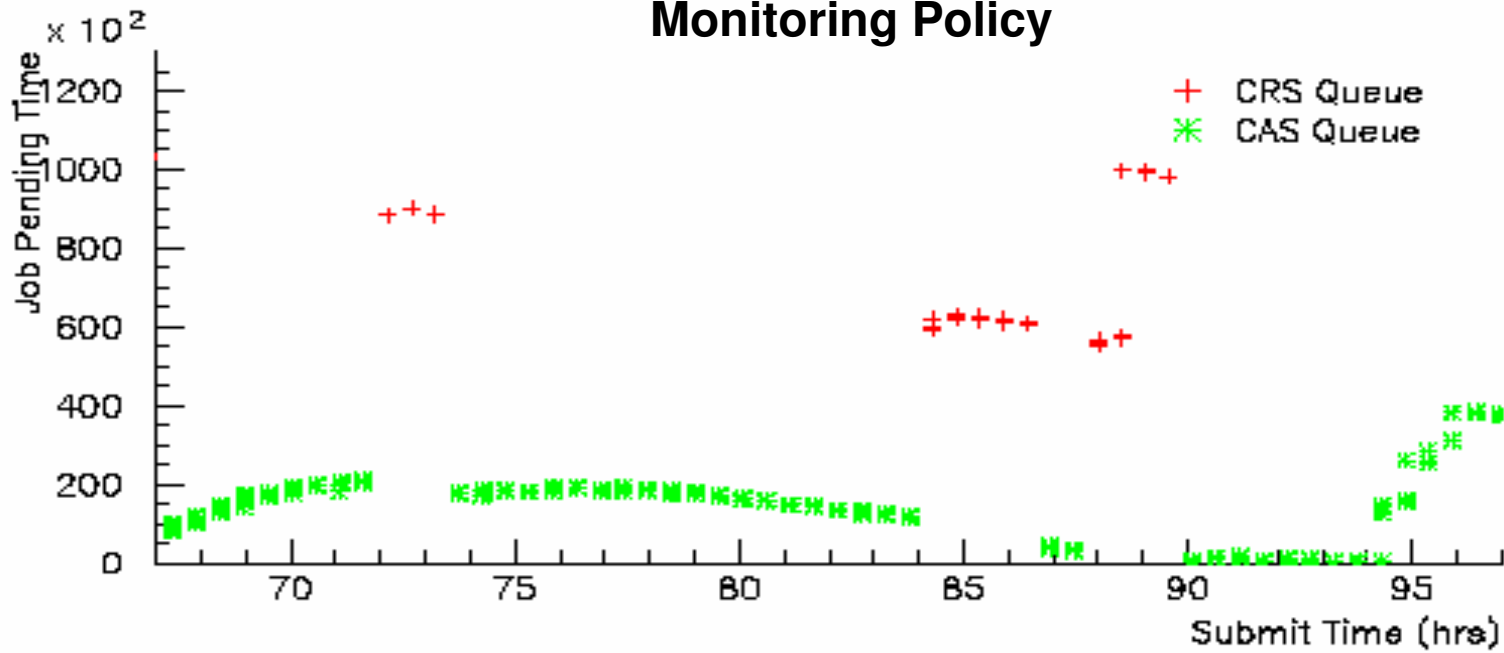
- In addition to providing a working scheduler for its users. SUMS is also used as a research tool mostly (but not exclusively) for grid research.
- Currently ongoing research projects include:
 1. Job scheduling with load monitoring tools
 2. Request Definition Language (RDL)

Job scheduling with load monitoring tools

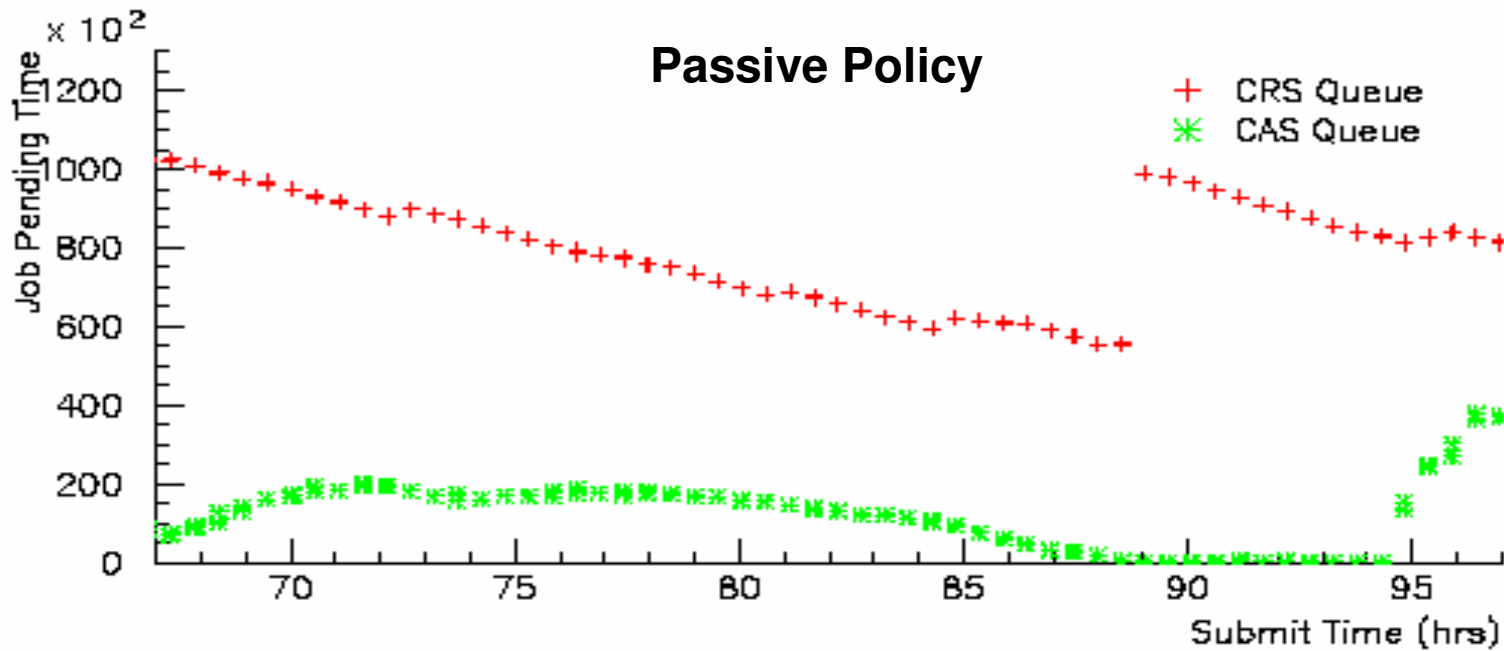
- The aim of this research is to devise simple, and efficient algorithms taking into account many variables for determining the best location to submit a job at the current moment in time.



Monitoring Policy


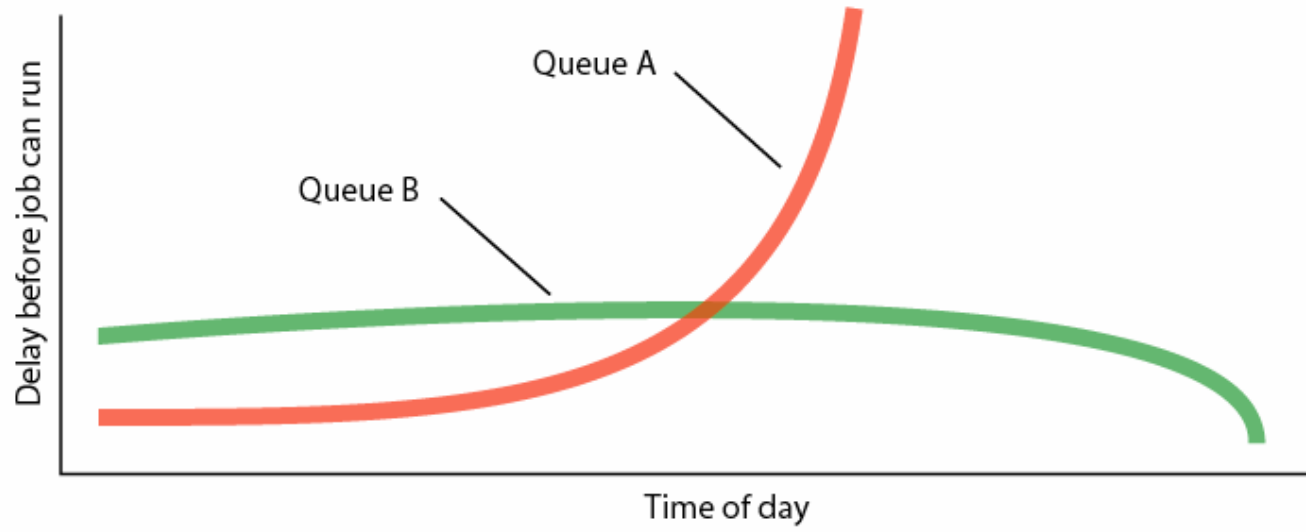


Passive Policy



Job scheduling with load monitoring tools

- Later test done during the week showed results not as good.
- Analysis shows fluctuating queue loads cause problems for this algorithm as it samples, and averages over 3 hours.
- **Observation:** It's inefficient to make a computation as intensive as the actual jobs your trying to submit.



Job scheduling with load monitoring tools conclusions

- A different equation is needed taking into account the differential (rate of change) of load and projecting forward in time.
- More testing will be needed.
- This will prove to be a valuable way of determining “where to submit”. The future looks good.



RDL (Request Description Language)

- **What if there existed a language so powerful that it could define not just a set of jobs but relationships between them not just NP or HEP but other fields as well.**
 - **Ex: Run these two jobs, take the one that's done first, merge the output run another process on it and draw me up some graphs.**
- **RDL is a new XML based definition language being developed by STAR in conjunction with Tech-X corp. that will allow for this level of abstraction and much more.**
- **Allows for Web-Service submitting and control**
- **RDL incorporates reality of every day work from different experiments.**
 - **Ex. J-Lab**

RDL v.s. JDL

RDL JDL

- **Submitting on a grid landscape**
- **Supports submit of multiple jobs**
- **Supports submit of multiple request**
- **Separates task and application**
- **Supports work flow**
- **XML format**

<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Who contributes to SUMS research and development

- PPDG – funding
- Jerome Lauret and Levente Hajdu – coding and administration of SUM at BNL
- Lidia Didenko – Testing for grid readiness
- David Alexander and Paul Hamill (Tech-X corp) - RDL deployment and prototype client and web service
- Eric Hjort, Iwona Sakrejda, Doug Olson – administration of SUMS at PDSF
- Efstratios Efstathiadis – Queue monitoring, research
- Valeri Fine – Grid testing
- Andrey Y. Shevel - administration of SUM at Stony Brook University
- And Others

