



CRAB3 and GlideinWMS



- Overview
- WMAgent Changes
- Glidein Usage
- Transition to CRAB3

Eric Vaandering

Glidein Workshop

UCSD

16 November 2011



CRAB3 is not CRAB2++



- What is CRAB3 and why are we doing it?
 - CRAB3 is a complete rewrite of CRAB. Only the name remains the same. No code in common.
 - Any preconceptions of “CRAB does X” w.r.t. Glideins should be discarded. It likely won’t do that anymore.
 - Uses WMAgents submission model, which I hope Matt just covered
 - Same client/server model as CRABServer
 - Much thinner client, all work done on the server
 - CRAB2 did all job preparation on the client



CRAB3: Motivation



- More stable development model
 - Based on WMAgent, the current CMS workflow software
 - Consolidated development, WMAgent designed for data
 - Modular structure allows us to add features we've said "later" to for years



Asynchronous Stageout



- In CRAB2, two ways for users to get output
 - Return through Condor “output_files_to_transfer” or gLite sandbox
 - srmcp from worker node to home SE at end of job
 - Both methods are problematic
 - In fact, biggest source of unforced errors in CRAB2
- CRAB3 scraps all of this and uses centralized FTS transfers
 - Job on WN writes to /store/temp/user
 - FTS jobs submitted to move files to home SE
 - Should increase CPU efficiency on WN, make transfers more reliable



Changes to WMAgent



- We've had to modify/extend some portions of WMAgent for CRAB3
 - Mostly related to proxies. User's proxy must be delegated to the worker node for writing files to local SE, srmcp of logs, and glxec under Glideins
 - gLite submission added to WMAgent
 - BossAir instead of BossLite (again, all new code)
 - I expect this will never be used in production
 - AsyncStageout is specific for users too
- What is not done yet is glxec switching on the submit node
 - High priority to get this done this year
- I expect Matt to have covered multi-site submission



Other WN Implications



- Input (user code) is handled differently than in CRAB2.
 - Stored on a central server, job (HTTP) requests from WN
 - Currently does not use squids/caching. Should change
- `/store/temp/user` (and maybe `/store/temp/group`) must be writable locally
 - Would be great to have Analysis factory pilots check for this



Open Questions



- As I said yesterday, we still don't have a great solution for discovering releases.
 - Currently assume the minimum: every site has every production release
 - This neglects sites that choose to have pre- or ancient releases
- Can we have releases advertised and put into the match condition?
 - Can WMAgent handle this?
 - Can we intelligently communicate back to the user if no match is possible?



CRAB3: Current Status



- Under active development. Early “alpha” versions in the hands of integration
 - Underlying WMAgent is used for all Tier1 work, not for MC and RelVals yet
- CRAB3 should be useful for some real workflows by expert users later this year
 - Remaining blockers are ability to restrict jobs to specific lumi sections and publication of results
- FTS based asynchronous stageout is written but not integrated yet
- Very little effort yet on modifications for local submission mode



Local Mode Plans



- One missing part of WMAgent is support for local schedulers
 - Scheduler plugins are easy, also need support for user switching with glxexec because of server architecture
 - Of course, this is ~same as the switching needed for Glidein submission
 - Delegated through myproxy
 - This is a requirement for the FNAL LPC Tier3
 - One of or the biggest analysis resources in CMS
 - Local Condor scheduling
- Bad news: Local mode will involve running a server



Comments



- Keep in mind that even though Glidein may do some things we need
 - We are not sure of the status of gLite with WMAgent
 - We need to allow support LSF, Condor, PBS, ARC
 - So, we need to aim for the lowest common denominator



Analysis Operations



- These plans are still being fleshed out
- Support model could change drastically with CRAB3
 - CRAB3 WMAgent is ~the same as production, so perhaps operators will be the same and AnOps does not really need expertise running a CRABServer
 - Some parts of CRABServer (ReqMgr, Stageout, etc) will be centralized at CERN.
 - Perhaps UCSD will run a WMAgent that pulls jobs from the central server
 - Unclear if we need as many WMAgent instances as we have CRABServers now
- Perhaps AnOps concentrates more on user support



Transition to CRAB3



- 2012 will be a year of transition
- Confident CRAB3 will be more reliable than CRAB2
- Expect that as feature set of CRAB3 grows, more people will move over
 - Can take over some non-CRAB2 workflows as well. (e.g. FWLite workflows)
- Still expect that CRAB2 will be supported through most of 2012
 - CRAB2 on the FNAL LPC will be supported until a replacement is available