SCAS Status

Jeff Templon, Nikhef

NIKHEF pdp

Memory Leaks

* SCAS is one process to many connections, whereas eg gatekeeper is single process per connection ... leaks accumulate across connections
* being fixed as we find them
* "refresh" is the workaround (a la apache)

refresh a la apache

* SCAS "master" process is just there. No restart

* internally, SCAS forks a slave process to handle connections. This is restarted every 300 seconds, internally by the SCAS master. No "service restart".

NIKHEF pdp

Errors

* Restart of slave each 300 sec should lead to a low-level error rate of approx 0.02%

* Should have realized this was not the explanation when SA3/CERN measured 0.2%.

* Problem : IO library released in wrong order during refresh - fixed.

Error rate now

* Upper limit 0.001% based on tests 6 March (if I compute right)
* The raw error rate of 0.02% is probably there, but is handled by client retries with improved strategy.
* This is with 10 worker nodes. Multiple SCAS servers can be used (ready for test) in case we hit a limit.

Future

* improvements in globus lib:
  * memory less leaky
  * bind to socket outside XACML lib: true Apache behavior, prelim tests show perf. scales with nr. of slaves
* production tests of multiple SCAS at Nikhef
* Oscar goes on holiday; multiple people know code.

NIKHEF pdp