Overview of service management issues

A. Forti

EGEE '07

Budapest, 5 October 2007

www.eu-egee.org

Information Society and Media

Talk Layout

- **Introduction**
- **Some problems**
  - Load balancing and high availability
  - WMS/LB
  - …..
  - Storage and file catalogs
- **Logging**
- **Documentation**
- **Who?**
- **Conclusions**
- **Discussion**

Introduction

- **June Operations Workshop in Stockholm presentation of the admin perspective of a release update cycle.**
  - http://indico.cern.ch/contributionDisplay.py?contribId=25&confId=12807
- **The purpouse of that talk was to provide some specific example of system administration operative problems and possible solutions to discuss with SA3 and within SA1.**
  - There was a follow up with SA3 in which many points were agreed. See SA1 tasks in savannah.
- **The purpose of this talk is the similar.**
- **Sys admin perspective with software problems normally translates in tools requests.**
  - I cannot do this please provide a tool.
- **So mostly this is a talk about administration tools that are not there.**
  - BUT hopefully also about solutions to avoid the need of these tools

Load balancing and high availability

- **There is a desperate need for implementing load balancing and high availability mechanisms.**
- **Top and ancient request is the ability to configure UI and WN and RB/WMS to point to more than one top level BDII.**
- **Multiple RBs on UIs are now possible but not always configured at sites.**
- **This is a problem not only for system administrator but also for users and creates an amount of wasted time on each side.**
  - Help!!! I can’t find my data anymore at your site
  - Have you tried another BDII?
  - How do I do that?
  - Reconfigure your UI with a new BDII that’s the easiest….
  - How do I do that?
  - Contact your local sys admin.

Draining jobs on WMS

- **Problem draining the jobs for maintenance**
- **To close a WMS the sys admin has to make disappear the service**
  - reconfigure the WMS to point to non existent BDII or
  - restart it on a different port
- **In each case the WMS returns an error to the users instead of a 'The system is draining/closed'**
  - A better and cleaner way to drain jobs. Is it possible?
- **Tool to allow a simple way to open/close submissions for draining jobs, maintaining the possibility to get-status and output;**
  - Even if a better and cleaner way to drain jobs is possible a tool to do it is still necessary.

WMS monitoring

- **General principles in the following slides apply to other services.**
- **Tool to monitor the information supermarket status**
  - The BDII cache on the WMS is sometimes not refreshed properly and is either corrupted or empty.
  - Again something easy to check, but it would be easier if there was an automated alert.
- **Tool to remotely check the services status and the job load of the WMS/LB in a graphical way**
  - A WEB based tool actually exists
  - Produced by someone at CERN
  - Installed at least at CERN and CNAF
  - Heavily used by WMS administrators and VOs
  - Not distributed, officially maintained or documented.
    - Corridor voices said it was distributed at coffee time!
- **Need to collect what is already out there!**
  - Where did I hear this already?

WMS monitoring

- **Another important general point made**
  - "What you really want is some way to be able to find the information at all"
  - Indeed, sometimes is all a sys admins really want so they can write their tool without spending days 'tailing' different log files to put everything together (when they manage to do it anyway).
- **Tools to easily extract the logging info and status of users job**
  - This is a tool that give an admin, logged into the WMS, a complete overview of what's happening with the job, correlating automatically the information in all the various scattered log files.
- **The above request could be easily generalised to**
  - Tools to easily extract information from log files to give an admin logged in node nodetype a complete overview of what is happening with the job/user/file correlating automatically the information in all the various scattered log files.

Garbage collection WMS&LB case

- **Define garbage collection criteria for WMS/LB;**
  - Apart from the 90 days there is no rule about what needs to be kept and what can be binned.
    - Is it only a sys admin problem?
- **Even in the case a sys admin decides autarchically what is garbage and what is not**
  - Tools for cleaning sandboxes, old condor files and old logs
  - Tools for cleaning the LB database
- **These tools could be avoided with a better behavior of the system in terms of cleaning its own garbage.**
  - Even in this case what becomes garbage and what not should be configurable.
- **All of the above is true for other services. Side effects for a sys admin are**
  - Space management problems
  - Performance degradation when databases are involved
  - Increased backup problems
  - General difficulties to find any information needed when required just for the shear amount of information to parse.

Uniformity of interfaces

- **Problem: add/ban a users/voms groups ;**
  - It is true that with LCAS it is simple to edit ban_users.db and a tool for that is superfluous
  - It is also true that at the moment not all the services use LCAS but WMS is not at the moment maybe it will change in the future.
- **Is this the only instance of things done differently on different services?**
  - Of course not.
  - And “of course editing low level files allows you to do everything, but you are tied to implementation that can change in time and from component to component. **Editing files is exactly what I don't want to do all the time**. What I would like to have is an admin interface that masks most of the low level files to the administrator and can be created bundling together all the tools that we listed here. If it easy to ban a user well, it will be easier to implement that piece of code in the admin interface.”

VO management

- **Tools to add/remove/reconfiguring VOs**
  - Most sites use YAIM to add VOs, but it reconfigures all the nodes.
    - Some nodes you can bear to reconfigure but some others you really DON'T want to touch. SEs in particular.
  - More sophisticated sites use their local fabric management tool
    - They still have to fiddle around to find the right things to do.
- **There is no tool and the procedure for each node type is not documented.**
  - Even my favorite going through YAIM code in this case is tough,
    - Bits and pieces are spread all over.
  - Maybe better with new YAIM? Forgot to ask the other day.
  - Still we need a tool to do it.
- **Same tool can be used for configuring VOMS groups?**
  - Perhaps. Elaborate and discuss the statement.

Node management

- **start/stop/status/restart for all MW services in /opt/edg/etc/init.d or /opt/glite/etc/init.d;**
- **start/stop/status/restart for all node types in /etc/init.d - this is the most important point, and those scripts must be working**
  - Currently e.g. WMS/LB cannot be stopped successfully without manual killing of left-over processes after gLite is stopped!
- **Tools to manage batch systems adding/removing/draining/closing queues adding/removing nodes and scheduling configuration.**
- **Tools to change core services on the nodes (BDII and MON) without reconfiguring the whole node**
  - Without reconfiguring the whole node is the key word for me.
  - But also to change the BDII is load balancing cannot be implemented
- **Tools to configure the firewall with the appropriate set of ports for the node type or service.**

Storage management

- **Tools to give a clear view of the relationships between disks, pools, files on disk and files in the namespace.**
  - In dcache it is possible but not easy and requires again that particular knowledge of how to extract the information and put it together.
- **Tool to syncronise Disk – SE database.**
  - Sometimes files in pools are not in the SRM DB and viceversa.
- **Having the ability to easily reconfigure the available disk among different VOs.**
  - Manchester hasn't assigned any pool yet because once it is done it is not possible to go back or resize.
- **Ability to assign space to VO subgroups.**
  - Even if the SRM understands VO sub-groups lcg-utils don't because they see only at VO level. At the moment fake VOs are created on storage elements as a work around.
- **A tool that allows to add/reconfigure and remove all traces of a VO from a SE.**
  - I know, I dedicated a slide to this, but is top priority for sys admins whatever their specialization
- **A dCache-specific topic (is it?):**
  - Improving the log files.
  - The dCache people know about this, but it is not known when it is coming.

Storage&catalogs

- **Typical case an SE with THOUSANDS of files registered in different LFC's. For any reason the SE needs to be replaced and the same FQDN cannot be maintained or the files path needs to change as in a move from DPM to dcache (or viceversa).**
  - Painful process of finding and contacting all the users who have data on the SE, ask them to un-register the replicas at the site,
  - Transfer the files with other tools than lcg-utils
    - I used FTS to transfer >24000 files from one SE to another in the same room.
  - Ask the users to re-register the files again.
    - Actual transfer might take few days the whole operations takes few weeks.
- **This is a problem that might happen few times but when it happens is almost catastrophic.**
  - And it has happened 4 times already in NorthGrid and other sites keep on bringing it periodically up.
- **The outlined solution is possible now that the users are still numerable but known. It will not be possible in the future.**

Storage&catalogs

- **Even without ‘catastrophic’ events as described in the previous slide synchronization between catalogs and storage is a problem. So here is the shopping list**
  - Tool to check the consistency between the entries registered in a catalog and the files on the SEs;
  - Tool for bulk removal of files and entries on storage and catalog if a SE died or if there are consistency problem
    - Bulk removal of files and entries is useful also for periodic cleanups when full productions become obsolete.
  - Define garbage collection criteria for files on SE not registered in catalog;
    - Same as the WMS case.
  - Implementing correct matchmaking for SEs in downtime or when the SE GIIS is down (machine crashed or something).
    - SE in downtime is almost there needs the final touches.
    - I’m not sure what happens when the SE GIIS is down.

Logging some input for later

From https://twiki.cern.ch/twiki/pub/EGEE/EGEEgLite/logging.html

- **“Logging serves multiple purposes”**
  - Number 1 logging reason for sysadmin : to diagnose problems! They do not check that everything is alright looking at the logs.
    - *Logging should be minimal in case of proper functioning, but should become much more verbose when an abnormal condition is detected.*
- **“It would be good to have only one logger library for each programming language and where the preferred library is the "standard" one rather than the newest”.**
  - YES!!!!
- **“Don't do too much logging“**
  - If the logging level is configurable and messages can be turned ON and OFF this is not a problem.
  - Don’t print messages without information
  - Use a single timestamp format easy to sort and grep
    - For example yyyy-mm-dd : hh:mm:ss

Logging (2)

- **Uniform scheme for log files location**
  - RGMA is a good example ls /var/log/glite
    - rgma-gin.log **rgma-server/** S20rgma-servicetool.log
    - …
- **“The overall programming approaches and coding conventions are diverse, and it is felt to be unproductive to require strong unification in the approach to logging.”**
  - Let’s discuss this stement. ☺
  - It is the first and last line any sys admin (apart from Jeff and I probably) has read scanning the document.
    - I can see their thoughts: “Right…..”
- **This document must be followed by the developers starting from yesterday!!!**

Documentation

- **User and Reference guides**
  - Command options and API description
- **Administration guide**
  - DPM has good system administration documentation
    - But many other services don't have an administration guide.
- **The full set of services which should be running for each node type is not documented.**
  - What is relevant for a sys admin
  - What is not relevant
  - Description of configuration files for each daemon
  - Description of what is in the log files
  - Where are the log files anyway?

- **Who is going to write it?**

Who?

- **Few times along the talk I asked 'Who's going to do it?'**
  - We don't want only to make a wish list, but we are available to collaborate with JRA1 for the implementation of all this things.
- **On the model of the WN WG maybe we can create a WG for each service or node type composed by at least**
  - sys admin
  - main software developer
  - yaim developer?
- **To make sure that the service or node type has**
  - Administration tools
  - Adequate documentation
- **Is this something that could be discussed?**

Conclusions

- **Basic system administration is underdeveloped.**
- **Tools or scripts to solve some of the listed problems exist, but they need to be bundled in a single suite and possibly with a single access point.**
  - Home made and not distributed.
- **Sometimes these tools don't exist and need to be written and maintained and documented.**
- **Some problems could be corrected with a better implementation rather than with tools (for example garbage collection).**
  - While you are mopping, it is nice to know that someone is fixing the plumbing. Jeff
- **Some problems could be solved by sys admin if there was a better documentation**
  - Documentation is needed when the system is not self-evident.
- **Need cooperation between sys admins and developers**
  - Proposal of a working group per service that oversees the creation of tools and documentation.

Thanks

- **Many Thanks for their comments to:**
  - Daniele Cesini
  - Graeme Stewart
  - Greig Cowan
  - Jeff Templon
  - Mario David
  - Paolo Veronesi
  - Steve Traylen

Discussion

???