Selected Security Issues

Authentication, Audit Logging, and reflections

David Groep, NIKHEF

Overview

- Authentication
  - namespace constraints
  - certificate profile
- Audit Event Logging
  - near-term guidelines document
- Security complexity and some personal reflection
  - CA release procedure
  - glExec and site acceptance of security middleware
  - VO naming
  - Do ‘normal’ people still understand us?

International Grid Trust Federation

IGTF
International Grid Trust Federation
AP | EU | TAG

Federation of 3 Regional "PMAs", that define common guidelines and accredit credential-issuing authorities

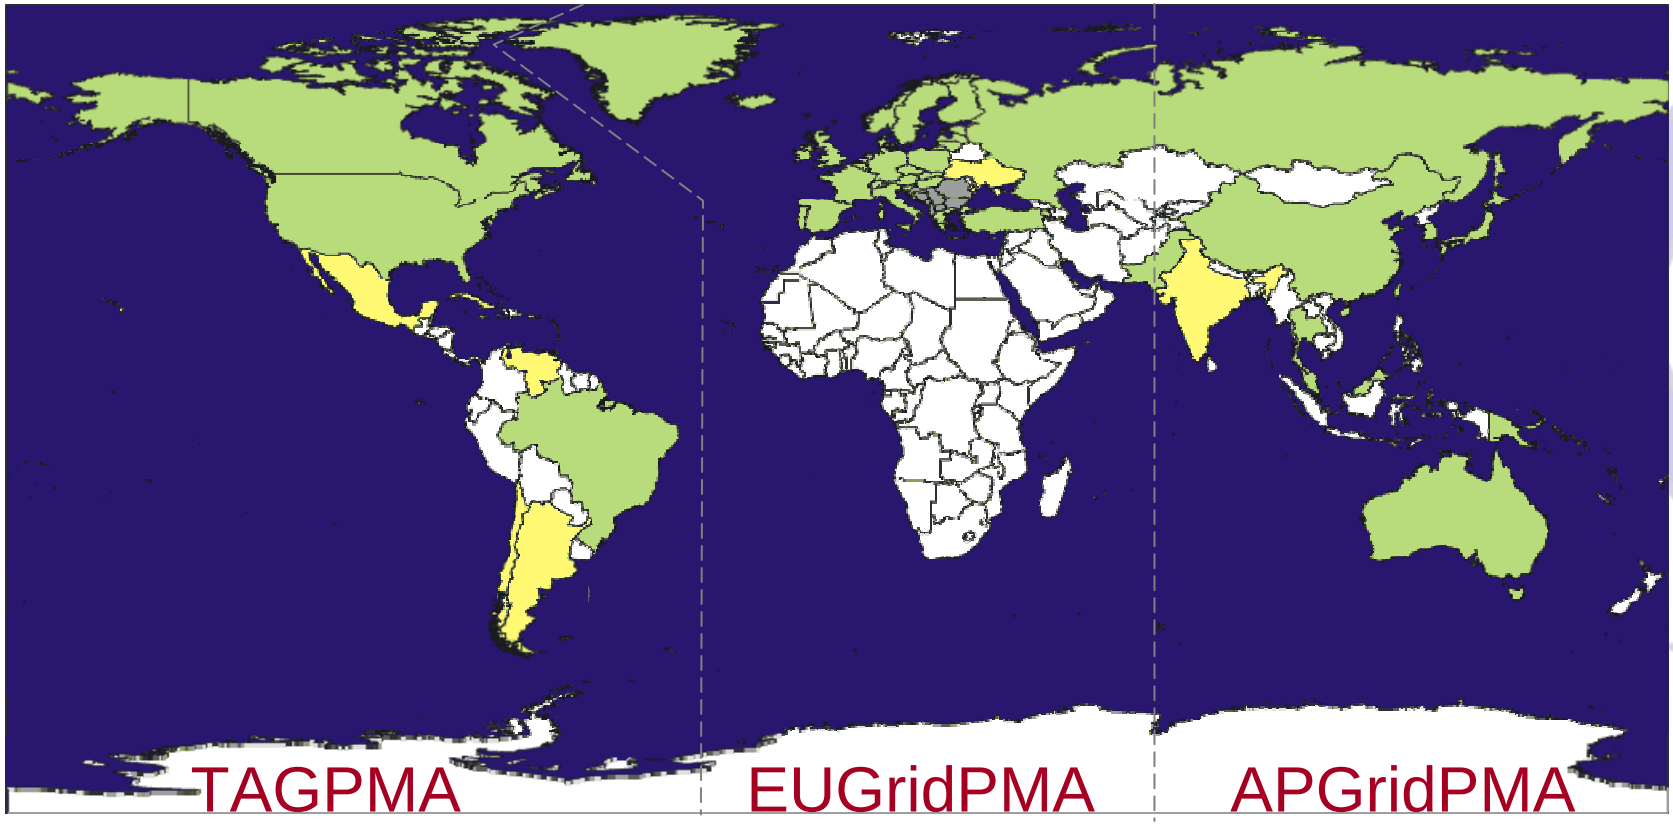

Relying Party issues to be addressed

Common Relying Party requests on the Authorities

1. standard accreditation profiles sufficient to assure **approximate parity**

   effectively, a single level of assurance sufficed then for relying parties – is changing today, as more diverse resources are being incorporated

2. monitor [] signing namespaces for **name overlaps**
3. a **forum** [to] participate and raise issues
4. [operation of] a **secure collection point** for information about CAs which you accredit
5. **common practices** where possible

list courtesy of the Open Science Grid

Guidelines: common elements in the IGTF

- Coordinated namespace
  - Subject names refer to a unique entity (person, host)
  - Usable as a basis for authorization decisions
  - This name uniqueness is essential for *all authentication profiles*!
- Common Naming
  - Coordinated distribution for all trust anchors in the federation
  - Trusted, redundant, sources for download, verifiable via TACAR
- Concerns and 'incident' handling
  - Guaranteed point of contact
  - Forum to raise issues and concerns
- Requirement for documentation of processes
  - Detailed policy and practice statement
  - Auditing by federation peers

Authentication - Namespaces

- Coordinated namespace can and must be enforced at the Relying Party end
  - the IGTF states explicit ‘domains of acceptance’ based on the subject namespace
  - especially important for hierarchies (such as the old SwissSign)

- GT1.0+ provided the ‘signing_policy.conf’ concept
  - shipped by the IGTF for each CA distributed
  - even in two formats

  - GT3.x+ Java part ignores this
  - only *very* recent versions of gLite Trust Manager are OK
  - status of other middleware is unknown

- there is no alternative mechanism – please implement and honour (and RPs ought to pressure their M/W providers)

signing policy requirements

- Support dynamic hierarchies*
  - describe namespace constraints for root and any and all subordinates in a single document
- also support *exclusion* of sub-trees
  - tackle the SwissSign issue

- See corresponding OGF CAOPS-WG draft document
  - and help by contributing to the doc and implementing it
  - implement either one of the currently distributed formats
  - or come up with a new one and work with the CAOPS-WG

** also needs OCSP, or path constructing validation service*

signing_policy.conf style

```
################################################################################
# Autogenerated EACL for 47d3d1a0 from switch-orgs.in# @(#)$Id$
#
# --- (temporary) workaround for http://savannah.cern.ch/bugs/?21033

 access_id_CA  X509    '/C=CH/O=Switch - Teleinformatikdienste fuer Lehre und Forschung/CN=SWITCH Personal CA'
 pos_rights    globus  CA:sign
 cond_subjects globus  '"/C=CH/O=ETH Swiss Federal Institute of Technology Zurich/*" "/C=CH/O=Fachhochschule Aargau Nordwestschweiz/*"
      "/C=CH/O=Paul-Scherrer-Institut (PSI)/*" "/C=CH/O=Universita della Svizzera Italiana/*" "/C=CH/O=Universitaet Basel/*"
      "/C=CH/O=Universitaet Bern/*" "/C=CH/O=Universitaet Luzern/*" "/C=CH/O=Universitaet St. Gallen/*" "/C=CH/O=Universitaet Zuerich/*"
      "/C=CH/O=Universite de Fribourg - Universitaet Freiburg/*" "/C=CH/O=Universite de Geneve/*" "/C=CH/O=Universite de Lausanne/*"
      "/C=CH/O=Universite de Neuchatel/*" "/C=CH/O=SWITCH - Teleinformatikdienste fuer Lehre und Forschung/*" "/C=CH/O=Switch -
      Teleinformatikdienste fuer Lehre und Forschung/*"'

# --- for /C=CH/O=Berner Fachhochschule - Haute ecole specialisee bernoise/*
access_id_CA  X509    '/C=CH/O=Switch - Teleinformatikdienste fuer Lehre und Forschung/CN=SWITCH Personal CA'
pos_rights    globus  CA:sign
cond_subjects globus  '"/C=CH/O=Berner Fachhochschule - Haute ecole specialisee bernoise/*"'

# --- for /C=CH/O=ETH Swiss Federal Institute of Technology Zurich/*
access_id_CA  X509    '/C=CH/O=Switch - Teleinformatikdienste fuer Lehre und Forschung/CN=SWITCH Personal CA'
pos_rights    globus  CA:sign
cond_subjects globus  '"/C=CH/O=ETH Swiss Federal Institute of Technology Zurich/*"'

# --- for /C=CH/O=FHS St. Gallen Hochschule fuer Angewandte Wissenschaften/*
access_id_CA  X509    '/C=CH/O=Switch - Teleinformatikdienste fuer Lehre und Forschung/CN=SWITCH Personal CA'
pos_rights    globus  CA:sign
cond_subjects globus  '"/C=CH/O=FHS St. Gallen Hochschule fuer Angewandte Wissenschaften/*"'

# --- for /C=CH/O=Fachhochschule Aargau Nordwestschweiz/*
access_id_CA  X509    '/C=CH/O=Switch - Teleinformatikdienste fuer Lehre und Forschung/CN=SWITCH Personal CA'
pos_rights    globus  CA:sign
cond_subjects globus  '"/C=CH/O=Fachhochschule Aargau Nordwestschweiz/*"'

# --- for /C=CH/O=Fachhochschule Solothurn Nordwestschweiz/*
access_id_CA  X509    '/C=CH/O=Switch - Teleinformatikdienste fuer Lehre und Forschung/CN=SWITCH Personal CA'
pos_rights    globus  CA:sign
cond_subjects globus  '"/C=CH/O=Fachhochschule Solothurn Nordwestschweiz/*"'

# --- for /C=CH/O=Fachhochschule Zentralschweiz/*
access_id_CA  X509    '/C=CH/O=Switch - Teleinformatikdienste fuer Lehre und Forschung/CN=SWITCH Personal CA'
pos_rights    globus  CA:sign
cond_subjects globus  '"/C=CH/O=Fachhochschule Zentralschweiz/*"'
```

8 files like this for the hierarchy …

.namespaces

```
################################################################################
#NAMESPACES-VERSION: 1.0
#
# @(#)$Id: 7b2d086c.namespaces,v 1.3 2007/02/08 14:24:44 pmacvsdg Exp $
# CA alias    : SWITCH hierarchy from SwissSign-Root (G1)
#    subjectDN: /C=CH/O=SwissSign/CN=SwissSign CA (RSA IK May 6 1999 18:00:58)/Email=ca@SwissSign.com
#    hash     : 7b2d086c
#
TO Issuer "/C=CH/O=SwissSign/CN=SwissSign CA (RSA IK May 6 1999 18:00:58)/emailAddress=ca@SwissSign.com" \
  PERMIT Subject "/CN=SwissSign Bronze CA/emailAddress=bronze@swisssign.com/O=SwissSign/C=CH"

TO Issuer "/CN=SwissSign Bronze CA/emailAddress=bronze@swisssign.com/O=SwissSign/C=CH" \
  PERMIT Subject "/CN=SwissSign Silver CA/emailAddress=silver@swisssign.com/O=SwissSign/C=CH"

TO Issuer "/CN=SwissSign Silver CA/emailAddress=silver@swisssign.com/O=SwissSign/C=CH" \
  PERMIT Subject "/CN=SWITCH CA/emailAddress=switch.ca@switch.ch/O=Switch - Teleinformatikdienste fuer Lehre
und Forschung/C=CH"

# SWITCH CA signs a Personal and a Server CA

TO Issuer "/CN=SWITCH CA/emailAddress=switch.ca@switch.ch/O=Switch - Teleinformatikdienste fuer Lehre und
Forschung/C=CH" \
  PERMIT Subject "/C=CH/O=SWITCH - Teleinformatikdienste fuer Lehre und Forschung/CN=SWITCH Personal
CA/emailAddress=switch.personal.ca@switch.ch"

# for SWITCH Personal and Server CAs

TO Issuer "/C=CH/O=Switch - Teleinformatikdienste fuer Lehre und Forschung/CN=SWITCH Personal CA" \
  DENY Subject "/C=CH/O=CERN/.*"

TO Issuer "/C=CH/O=Switch - Teleinformatikdienste fuer Lehre und Forschung/CN=SWITCH Personal CA" \
  PERMIT Subject "/C=CH/O=.*"
```

Certificate Profile

- Conversely, CAs are trying to cater for current M/W
  - abstain from using ill-supported subject DN elements (such as uid/userid, serialNumber)
  - restricted ordering of RDN components
  - encode policy oids in EECs *
    we can and should discuss how RPs could make use of that
  - when there is an installed M/W base, CAs will be willing to setup OCSP and/or hook into GTS and like services (LBNL/ESnet is running it as a pilot service)

Grid Certificate Profile

- Your input is welcome in the IGTF and CAOPS-WG
  - See GridForge for the latest Grid Certificate Profile document

- Also provide input on known ‘incompatibilities’
  - such as the use of dashes (‘-’) in host certificates
  - … review and comment into this (living) document

Security Auditing Logging

improving auditing and traceability capabilities in the short term

version 0.7+

David Groep

NIKHEF

Outline

- Rationale
- Message content
  - identification
  - time stamping
  - length and parsing
- What to log
  - server startup and termination
  - connection establishment
  - request logging
  - message linking and interrelation
  - aggregation
  - log completeness and accuracy

Rationale

Primary aims

- facilitating *post-factum* security auditing
- at the single-site level
- integrated in the existing site audit infrastructure
- be parsed automatically by software

Secondary aims

- prepare ground for future distributed aggregation systems
- help in trouble shooting is a nice side-effect

Which M/W should log

The guidelines apply to

- all services
- all components that run
  under dedicated or privileged accounts,
- all components that are started and act
  on behalf of a user with enhanced privileges

Syslog

Security logging is to be integrated in site logging systems, so

- use of the syslog(3) facility MUST be supported
- and MUST be enabled by default

- middleware MAY also log to other destinations
- and MAY use indirect systems such as *log4cpp* &c

(guidelines are written in normative, RFC2119 language)

Syslog semantics reminder

Facility

- classification of events used to discriminate sources
- typically different *facilities* end up if different log files
- a few well known ones (DAEMON, FTP, USER, LOCAL1-7)

Severity ('priority')

- is the 'log level' and is strictly incremental
- from DEBUG, INFO, NOTICE, WARN, … up to EMERG
- is orthogonal to the set of facilities
- highest levels are really reserved for system emergencies

Message content: facilities

- the syslog(3) facility MUST be configurable
- messages that *could* contain re-usable private data
  MUST be logged to AUTHPRIV by default
- all other messages SHOULD be
  logged to DAEMON by default

Re-usable private data can appear in many unlikely places

For example,

it is well known that users tend to type a password when asked for the user name – that's why login(8) and ssh(8) use AUTHPRIV also for logging the username

Message content: identification

Messages from a single request, session MUST be relatable

- A unique identifier MUST be prepended to every message

- the identifier
  - MUST contain the (shortened) process name ‘in.ftpd’ (“*TAG*”)
  - MUST contain the process ID between square brackets ‘[4213]’
  - a single colon plus space (“: ”) character MUST follow process id

- if a single process handles more than one request
  - a unique *session identifier* MUST be generated and appended after the colon-cum-space

Message content: time stamp

- normal time stamping is taken care of by syslog

- *if independent (or sub-second resolution) TS's are required*
  - the timestamp MUST appear as the first item after the *msgid*
  - the timestamp MUST be in the ISO8601 '*combination of date and time of day representation in the extended format*'
    like: `2006-11-15T10:13:02Z`
  - MAY contain fractional second information as per ISO8601
  - MUST include the time zone, and zone either
    - MUST be in UTC, indicated by "Z"
    - or MUST be expressed numerically "+01:00" or "+0100" (*only log4j*) since abbreviations like "PST" are not unique globally
  - date and time specification MUST NOT contain whitespace
  - timestamp SHOULD be logged in UTC ("Zulu time")

Message content: length, parsing

- Note that length of a syslog message on many systems is
  - 1024 octets max (some implementations can eat only 1024!)
  - including the system-prepended time and host information *(typically up to 64-80 characters)*

- Messages MUST be machine readable
  - so use of white space MUST be consistent within each service

- use of "*name=value*" pairs is RECOMMENDED
  - if used, *name* and *value* MUST be separated by a single equal ("=") sign
  - a *value* containing white space MUST be quoted in double-quotes (“”)
  - the *name* SHOULD consist solely of the characters "a"-"z", "A"-"Z", "0"-"9" and "_"
  - if a *value* contains a double-quote (“”) character, it MUST be escaped with a single backslash ("\")

Message characters

Allowable characters:
MUST be seven-bit ASCII in an eight-bit field. In this code set, the only allowable characters are the
visible characters (%d33-126) and space (SP value %d32).

However, no indication of the code set used within the MSG is required, nor is it expected. Other code sets MAY be used as long as the characters used in the MSG are exclusively visible characters and spaces similar to those described above. The selection of a code set used in the MSG part SHOULD be made with thoughts of the intended receiver. A message containing characters in a code set that cannot be viewed or understood by a recipient will yield no information of value to an operator or administrator looking at it. BUT we want to be understood widely

from RFC3164

Examples

- Examples in the document (non-normative) appendices

```
daemon:notice
   jss-serv[5241]: event=NewConnection
        ts=2006-09-28T10:09:23.021Z
        remoteHost=192.16.199.115:28773
        DN="/DC=org/DC=example/CN=Pietje Puk"
daemon:debug  (can be debug severity since the pid remains the same)
   jss-serv[5241]: event=info
        msg="delegating authorization to mapper in-line"
        pid=5241
daemon:notice
   jss-serv[5241]: event=AuthenticationRequest
   msg="access allowed" DN="/DC=org/DC=example/CN=Pietje Puk"
        uid=43004 uname=dzero004 pgid=2008
        gname=dzero sgids=512,100
...
```

What to log?

Start-up, termination

- Log on
  - start-up
  - termination(!)
  - reconfiguration
- at severity NOTICE
- with information
  - configuration used or (re)loaded – if service can be configured
  - on termination: reason of termination or status (success) code
- if termination is due to a signal (except KILL)
  - log severity MUST be raised to WARNING

Service access

1. On each access the source address and port – or the process id and owner of the process connecting to the socket – MUST be logged at severity **LOG_NOTICE**.

2. If the connection is TLS-authenticated with client credentials, the certificate subject DN of the client MUST be logged at severity **LOG_NOTICE**. If the TLS handshake fails, a message including the reason, and if possible the client subject name, MUST be logged at severity **LOG_WARNING**, unless the handshake failed because no data at all was received from the remote peer, when the severity MAY be set lower but not lower than **LOG_NOTICE**. This message MUST re-iterate the connection source information reported in (1).

3. The result of the (subsequent) authorization decision that results in granting access (i.e. the allow decision) MUST be logged at severity **LOG_NOTICE**. If a local identity is assigned to this session or request, the local mapping, including the numeric *uid*, the numeric *gid*, and the list of numeric supplementary *gid*s, MUST be logged at severity **LOG_NOTICE**.

4. The result of a denied authorization MUST be logged at severity **LOG_WARNING**. This message MUST re-iterate the connection source already reported in (1)&(2).

5. In either case, the list of attributes (e.g. obtained from VOMS attribute certificates), or the fact recognising the lack of any attributes, MUST be logged at **LOG_INFO**.

6. Any (session) identifiers that are required to link the authorization state to subsequent specific requests MUST be logged at severity **LOG_NOTICE**.

Service access

Other information at this stage MUST NOT be logged with a severity higher than LOG_DEBUG

- thus, if you want to add more diversity for debugging problems, you MUST do this in a debug-verbosity system that's orthogonal to the syslog severities
- if you want to boost performance, suppress calls to syslog for DEBUG messages unless a service specific debugging flag is on

---

At the end of a request

- If the service supports multiple requests with a single access, and if the termination of such an established session can be identified, a message at severity LOG_NOTICE SHOULD be logged at completion of each session
- Such a message MUST contain the unique identifier that ensures this message can be linked to the initial access messages

Request logging

- For external requests that successfully modify persistent state for which the service is responsible, both the action and the operands MUST be logged at severity LOG_INFO.
  If such an external request fails, it SHOULD be logged at any severity level that is considered appropriate, but not at a lower level than LOG_INFO.

- External requests that do not modify the persistent state and complete successfully MUST NOT be logged at a severity higher than LOG_DEBUG.
  External requests that do not modify the persistent state that failed MAY be logged at any severity level up to but not higher than LOG_INFO.

- Other (internal) actions MAY be logged at any severity up to but not exceeding LOG_INFO.

Linking messages

- A ‘linking’ identifier between a service request and any and all access log messages of a single service MUST be provided in all messages (e.g. using a session id).

- If a service is able (within reasonably bounds) to obtain an external request or session identifier, this identifier SHOULD be logged as part of at least one message that also contains the single service request identifier.

- If a service delegates (part of) the processing of a request or session to another service or process or thread, or if it contacts another service to satisfy parts of a request, a unique identifier referencing this delegation process or thread or service MUST be logged by the parent or invoking process as part of at least one message that also contains the single service request identifier.
  The completion of such a delegated request, or the termination of the process or thread, MUST be logged by the parent or invoking process.
  The message severity for all such messages MUST be set to LOG_NOTICE

Aggregation

- Multiple pieces of information MAY be combined into a single log message, as long as all information is to be logged at the same severity level. If the service only ever sends a single message to syslog, all information MAY be combined into a single log message, regardless of the severity level, and the severity level of this single message MUST be set to the highest severity of any of its constituent parts.

- If a service instance only ever serves a single request, it is RECOMMENDED that the service start up message and the connection establishment message are combined into a single message

How and when to log

In case a service contacts other services to fulfil a request, or initiates processes to which processing is delegated, it is

- strongly RECOMMENDED that log messages are generated and sent out even if the delegation or invocation fails, even if such a failed attempt will cause the service itself to terminate.

- Note that this could imply wrapping any such invocations in code that traps and recovers from exceptions to ensure the log message is generated and sent.
- *I would rate this as a lower-priority item if time-pressed*

Examples

- Examples in the document (non-normative) appendices

daemon:notice
jss-serv[5241]: event=NewConnection
ts=2006-09-28T10:09:23.021Z
remoteHost=192.16.199.115:28773
DN="/DC=org/DC=example/CN=Pietje Puk"

daemon:debug *(can be debug severity since the pid remains the same)*
jss-serv[5241]: event=info
msg="delegating authorization to mapper in-line"
pid=5241

daemon:notice
jss-serv[5241]: event=AuthenticationRequest
msg="access allowed" DN="/DC=org/DC=example/CN=Pietje Puk"
uid=43004 uname=dzero004 pgid=2008
gname=dzero sgids=512,100

...

Guidelines Document

- Discussions and comments to me or to the security audit mailing list

  security-audit-log-discuss@george.lbl.gov

  - *mailing list will also discuss future directions in context of the new SciDAC project on distributed grid auditing*

- The document in the SCG area of the EGEE-II EDMS tree

  https://edms.cern.ch/document/793208

- latest version: 0.7
  - has been circulated extensively, can we agree now?

My own ‘wish’ list

Some personal gripes on management issues of our security middleware

‘lcg’ CA distribution – Issuer names – site glExec acceptance – VO naming – <X/>*<*S>>AC?>*><M></L> ‘clarity’ -

Why?

‘When you talk security, I’m always lost after two lines …’

‘… there are too many entities, identities and identifiers in this one single sentence …’

‘… but I already did this for that other component, why is it different here?’

‘… oh, there was *yet another* file where I had to set this policy?!’

‘… I was stung by one of those angle brackets of yours …’

Security middleware is becoming too complex
for many sites to either understand or handle

CA distribution in EGEE

- Geared towards frequent automatic updates at the sites
  - completely decouples from S/W release
  - based around a single meta-RPM with dependencies
  - sites use RPM package managers to update
  - possible automatically via cron
  - all packages are digitally signed
  - SFT/SAM tests verify installed versions several times per day
    - trigger warnings and errors on mismatch, that escalate procedurally
  - looks ‘fine’ from the user/site end, but …

CA Release Procedure

Introduction

This page intends to describe a procedure for the annoucement of a new CA release. The annoucement should be done in a well defined order to ensure consistent deployment within LCG/EGEE and avoid mess in SFT.

Teams involved

The following teams (and people) are involved :

- IGTF (David Groep)
- EGEE deployment team (Laurence Field, Oliver Keeble, Louis Poncet) man-install-grid-support@cern.ch
- EGEE Integration team (Marian Zurek, Joachim Flammer, Alberto di Meglio) project-eu-egee-middleware-iteam@cern.ch
- SFT development team (Frédéric Schaer, Piotr Nyczyk, Judith Novak, Rafal Lichwala) project-eu-egee-sa1-sft-devel@cern.ch (and for informationnal purpose it would be good to CC : project-eu-egee-sa1-cic-on-duty@cern.ch)
- LCG Security Officer, project-lcg-security-officer@cern.ch

Procedure

1. New CA release is announced by David Groep via the EUGridPMA-Announce list to :
   1. the LCG Security Officer via project-lcg-security-officer@cern.ch
   2. SFT team via the sft-devel mailing list (project-eu-egee-sa1-sft-devel@cern.ch)
   3. EGEE middleware deployment team man-install-grid-support@cern.ch
   4. EGEE integration team project-eu-egee-middleware-iteam@cern.ch

   This announcement should not be done on LCG-ROLLOUT !

   Either the Security Officer or the Deployment team can declare this update as urgent
2. David Groep (acting as SA1 this time) will also build the entire lcg-CA package, and make that available on a special site only intended for INTERNAL use by the deployment process. This "special" distribution would contain:
   - the IGTF "classic" CAs
   - the "lcg-CA" meta package
   - the SwissSign hack ("ca_patch_eugridpma_gridppvuln_14013-1.0-1") package
   - the FNAL-KCA and any other LCG specific CAs

David will also test the full lcg-CA distribution.

These RPMs would be made available on a pretty obscure web site (not on production sites at that step)

3. David will create a GGUS ticket, with with a standar subject: "CA update, version X.Y.Z-R"

And with a comment inside saying: "please, assign this ticket to the SAME/SFT support unit". The integration team will also be involved in the ticket to start preparing the repository ("Involve others:" in GGUS), as both changes are independent. The ticket will contain
   - the URL, e.g.:
     - http://www.cern.ch/groep/cadist/lcgpreview-X.Y-Z/
   - the LCG specific change log (as an attachment)
4. SFT team modifies the CA test of SFTs, updates the CVS, and deploys it on SFT machine. Once the SFTs are updated, a timeout (7 + 1 days) starts that allows time for sites to upgrade. Sites are allowed 7 days to upgrade, while 1 extra day is added for this procedure to complete. Remaining time shall be printed in the SFT test.

   At that time, an announce email has to be sent to the same-announce list (same-announce@cern.ch), so that people running SFT/SAME instances know they have to upgrade their copies.

   In case the update is urgent, the timeout will be set to 1 day.
5. SFT team informs the Integration team that the SFTs are ready (this is done by re-assigning ticket to the integration team, project-eu-egee-middleware-iteam@cern.ch, and adding (mandatory) comments).
6. Integration team adds the new CA release to a temporary APT repository till they are tested and ready to be released. They send a mail to grid-cern-pps-admins@cern.ch to request the new CA to be tested, including the link to the temporary APT repository.
7. When the testing confirms that there is no problem, they move them to the production apt repository and update the CA related web pages (comments have to be inserted in the ticket). These pages currently are :
   - http://grid-deployment.web.cern.ch/grid-deployment/lcg2CAlist.html
   - http://grid-deployment.web.cern.ch/grid-deployment/download/RpmDir/security/index.html

   (proposal by DG: copy also the three files in the private preview directory to a proper location and make links to them from the lcg2CAlist.html page. These files are: the changelog announcement, the plain list of RPMs, and the Quattor template)
8. Integration team announces the new CA release using the **EGEE broadcast tool (select "To ROC Managers" and "To Production Site Admin")**, and **mailing (ccing) the glite-announce mailing list**. After this, they close the ticket. Template emails including email subjects can be found here. The changelog to be included in the mail is part of the GGUS ticket, and can also be found at http://www.cern.ch/groep/cadist/lcgpreview-X.Y-Z/ in the file `lcg-ca-X.Y-Z.txt`.

Completion of steps 5 to 8 is not expected to take longer than one day. Sites can then update from the standard apt/yum repository, or do that automatically from cron.

This procedure uses a GGUS ticket opened by DavidG (step 2), and ticket is closed by Integration team after step 7. GGUS ticket subject should be "CA update, version X.Y.Z-R".

To keep informed the LCG Security Officer, project-lcg-security-officer@cern.ch should be in copy of the ticket.

At any stage of the procedure, the LCG Security Officer will have the possibility for a veto. This could be done through the GGUS ticket.

No non-urgent CA upgrade should be started on fridays (because of deployment issues this could cause).

CA release process

- Defined as of IGTF distribution release 1.5
- For last 7 releases, 0 went smoothly
  - update announces before SAM tests were ready
  - internal inconsistency in the lcg-CA RPM
  - change log information lost in announcement
  - lcg-CA RPM does not match IGTF distribution
  - not the right people involved in the GGUS ticket
  - …
- this is an ‘infrequent’ and relatively well-controlled process!
- do we really want to repeat this for VOMS trust anchors?
- Or can we leverage innovations (like GTS) soon?

CA naming

- When a CA keypair changes, the name needs to change
  - the ".0"/".1" OpenSSL convention
    - not a good match for RPMs
    - too prone to mishaps when the ".0" file is gone
  - fetch-crl can retrieve only one CRL per CA hash
- so the CAs change name (or extend** lifetime)

- … leading to subsequent inconsistencies in VOMS
  - and the need to reregister for VOMS versions <= 1.6.x
  - *new version will have a switch to disable this!*

- *oh yeah … also and please fix the signing policy issues…*

glExec

- *a thin layer to change Unix credentials based on grid identity and attribute information*

- Deployment models
  1. a non-privileged dedicated CE or scheduler, accepting authenticated user jobs and submitting to the batch system
  2. on-demand CE, submitted by VO or user to a front-end system, that then receives user jobs and submits these to the batch system

  Novel models with pilot jobs

  – *glexec as a per-job declaration and limited sandboxing system*

  1. installed *setuid* and mandatory
  2. installed non-setuid, without actual enforcement (declaration-only)
  3. replaced by a null-operation (a simple ‘exec’ shell script)

glExec status

- Opinions on desirability of glexec *et al.* varies widely
  - FNAL has adopted and deployed it (as a first)
  - LCG T1 sites have mixed feelings
    - but an inadvertent side-effect of the FNAL deployment is that the *LHC experiments* are now pushing strongly …
  - smaller sites (most of EGEE and many T2s) are very, very reluctant

glExec – comments

- I don't want glexec on the worker nodes. It is a change of identity on the fly done by the application itself.
- It means that we trust VOs to know what they are doing because we don't have a certification procedure in place; it potentially allows VOs to circumvent local policies; it makes difficult to do accounting; […]
- With the place holder I can point the finger to one person, the one who runs the place holder; with a change of identity I cannot do that anymore. This is all assuming that
  - a) we have good faith on the other side or
  - b) that the system is not hacked.

  At that point anyone, whether authorised or not, on the grid can change identity on a system.
- We might as well give up control of our clusters.

verbatim comments from a site administrator, reflecting real concerns in the community!

glExec

- Are we again running ahead too fast?

VO naming

As agreed in March 2006 in the JSPG: DNS style naming

- The VO name is a string, used to represent the VO in all interactions with grid software, such as in expressions of policy and access rights.
- The VO name MUST be formatted as a sub-domain name as specified in RFC 1034 section 3.5. The VO Manager of a VO using a thus-formatted name MUST be entitled to the use of this name, when interpreted as a name in the Internet Domain Name System.
- This entitlement MUST stem either from a direct delegation of the corresponding name in the Domain Name System by an accredited registrar for the next-higher level sub-domain, or from a direct delegation of the equivalent name in the Domain Name System by ICANN, or from the consent of the administrative or operational contact of the next-higher equivalent sub-domain name for that VO name that itself is registered with such an accredited registrar.
- Considering that RFC1034 section 3.5 states that both upper case and lower case letters are allowed, but no significance is to be attached to the case, but that today the software handling VO names may still be case sensitive, all VO names MUST be entirely in lower case.

ATLAS atlas.ch

FUSION fusion.eu-egee.org

vlemed vlemed.vl-e.nl

gin.ggf.org gin.ggf.org (it does work!)

Some obstacles

- The YAIM configuration tool used the VO name in some internal definition variables (solved, IFAIK)
- VO specific software, traditionally stored in a directory referred to by

  `$VO_voname_SW_DIR`

  still uses the VO name in an environment variable
- This can easily be fixed, all tools are there, but requires coordinated change in more than one component
  - Will it still happen, at some future date?

Gap between development and deployment is growing

Deployment realities

- Many of the MW solutions are highly configurable
  - developers mostly see their own configuration
  - but software redistributors add their own configuration layer (e.g., YAIM; vdt will have something similar)
  - many of the configuration options are lost
  - and most remain *unset*
  - so the built-in default will be used
  - *so they must be at least fail-safe, but even better*
    do it right for majority of the real-life cases

Dilemmas?

- But then, what is ‘good’ and what is ‘bad’ is not even clear
  - some users see nothing wrong even with daemonized jobs
    - gets them free, non-accounted CPU time as well
  - what site admins care about depends on the phase of the moon
  - but still…

- Good documentation and site admin induction is at least as important as (new) middleware…

  - try to explain new middleware to your local sysadmin
  - avoid any and all stuff that mortals cannot understand
    - is X**ML always the most appropriate choice?
  - provide simple, easy to read and understand example policies
  - ship with safe but workable built-in defaults
  - **make it fit with what a typical (Unix) shop understands!**

Q&A ...