



Subversion in ATLAS

David Quarrie
Drquarrie@lbl.gov

History

- Detailed evaluation and proposal in 2005-2006 by Vincent Garonne & Christian Arnault (LAL, Orsay)
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/StudiesForUsingSVNInAtlas>
- Intention was to migrate in Spring-Summer 2006
- However, the migration to ROOT 5 and associated LCG-AA packages in Spring 2005 took a lot longer than expected and was very disruptive
- In Summer 2006 we re-evaluated the proposal to migrate to SVN and decided that we would have to postpone it until after ATLAS turn-on
 - Yet more disruption following ~6 months of disruption
 - Concern about closeness to ATLAS turn-on
 - Hindsight is wonderful - that got delayed by >12 months
- Next slides summarize the evaluation and proposal
 - Based on talk given by Vincent at ATLAS Software Week in Apr 2006



CVS *versus* SVN

- ▶ *SVN* really works faster than *CVS*:
 - Transmits less information through the network
 - Supports more operations for offline mode, e.g client sends deltas to server when committing

| Source code management system | CVS | SVN |
|---|-----|-----|
| Atomic Commits | No | Yes |
| Efficient (constant-time) tagging | No | Yes |
| File renames handles cleanly (renaming, moving) | No | Yes |
| Advanced Merging Features | No | Yes |
| Manipulation of all the file types | No | Yes |
| Easy Development on Branches | No | Yes |
| Easy Server Administration | No | Yes |
| Remote Repository Replication | No | Yes |
| Scalable Performance and Admin | No | Yes |
| Advanced repository permissions Features | No | Yes |
| Web Browser Interfaces | No | Yes |
| HTTP proxy to connect the server | No | Yes |

** For more infos, look [2]





Scenario for Atlas

Roadmap

1. Convert the Atlas CVS repository into a Subversion one
2. Provide tools to provide synchronization operations between CVS and SVN repositories
 - For a certain period, we will have both system running at the same time
 - remark: iterative transition is less painful for the developers :)
3. Incorporate SVN in the complex Atlas software management chain, it supposes impact to :
 - CMT
 - Tag Collector
 - Nicos
4. If successful, we could do the permanent switch; if not ... :(





Converting CVS Repository to Subversion

- ▶ We have selected the following organization for each project/package/tag:
`<project>/.../<package>/trunk/...`
`<project>/.../<package>/tags/<tag>/...`
- ▶ The complete Atlas cvs repository has been converted to a svn one. The svn url is :
`file:///afs/cern.ch/atlas/software/svnrepository/repository`
- ▶ It takes a looong time to do that the first time
 - Try different *SVN* backends: Berkley DB, SVN file system
 - It takes one week to complete the procedure, e.g. 8 hours for the AtlasCore project
 - But update will only concern recent tags



Conversion Time for 1 of 10 Projects





Conversion validation and management tools

Conversion validation tools

- ▶ We have implemented several tests to check the quality of the conversion:
 - Comparing the sources at each tag with diff tests
 - Comparing the list of tags for a package
 - Check global structure of the *SVN* repository
 - `/afs/cern.ch/atlas/scripts/testcvs2svn.py`

Management tools

- ▶ We have also implemented a script which re-organizes the *SVN* repository:
`<project>/.../<package>/trunk/...`
`<project>/.../<package>/<tag>/...`



Performance and scalability

Test stress scenario

We define the following **action** for an user:

- ▶ User selects randomly a package
- ▶ He extracts and modifies it, then registers it until the operation is successful
- ▶ Then he sleeps n seconds with n sort randomly $\in [1, T_{max}]$

Experiment parameters

- ▶ We have done experiments with different:
 - Total number of users (10, 100, 1000)
 - Total number of packages (10, 1000, 10.000)
- ▶ Concurrent user actions were launched by a multi-threaded program (`/afs/cern.ch/user/v/vgaronne/svnTestStress`)



Synchronization operations

Mainly at the user level:

- ▶ We proposed a script "asvn" that mimics the CVS syntax for SVN, including the synchronization operations on tags:
 - Non exhaustive commands are: checkin, checkout , tag, ...
 - Provides also cmt specific commands working with SVN, e.g. recursive checkout, svntags, svnsubprojects, ...
 - see [/afs/cern.ch/atlas/scripts/asvn](http://afs.cern.ch/atlas/scripts/asvn)
- ▶ We assume that when users switch completely to svn for a package, they will not work anymore with CVS for this package, otherwise we will have integrity problems
- ▶ The synchronization operation should be triggered at the user/developer level during tag operations
 - > `asvn --sync tag Package-00-00-01 Package`



Impact to the Atlas software management chain and tools

Impact to the work model

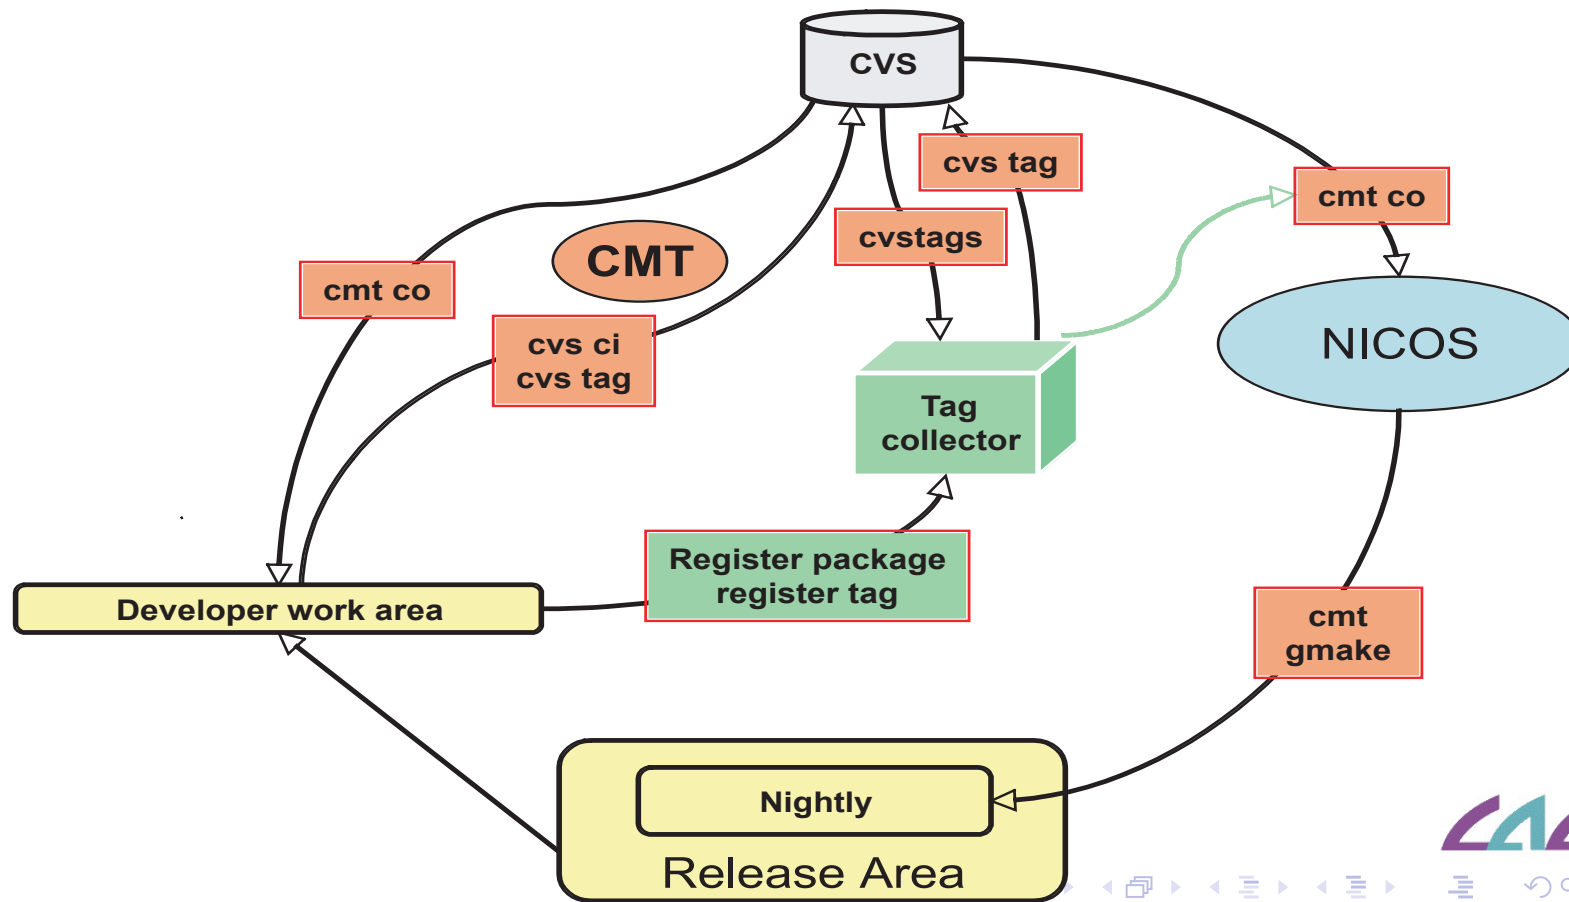
- ▶ No real change is expected

Impact to the tools

- ▶ CMT: The "asvn" script previously presented covers these aspects
- ▶ Tag Collector: Needs an interface to SVN, but should be much simpler than the interface to CVS. If not possible immediately, we could go on temporally with the synchronized Atlas CVS repository
- ▶ NICOS: Normaly only concerned with checkout operations ?

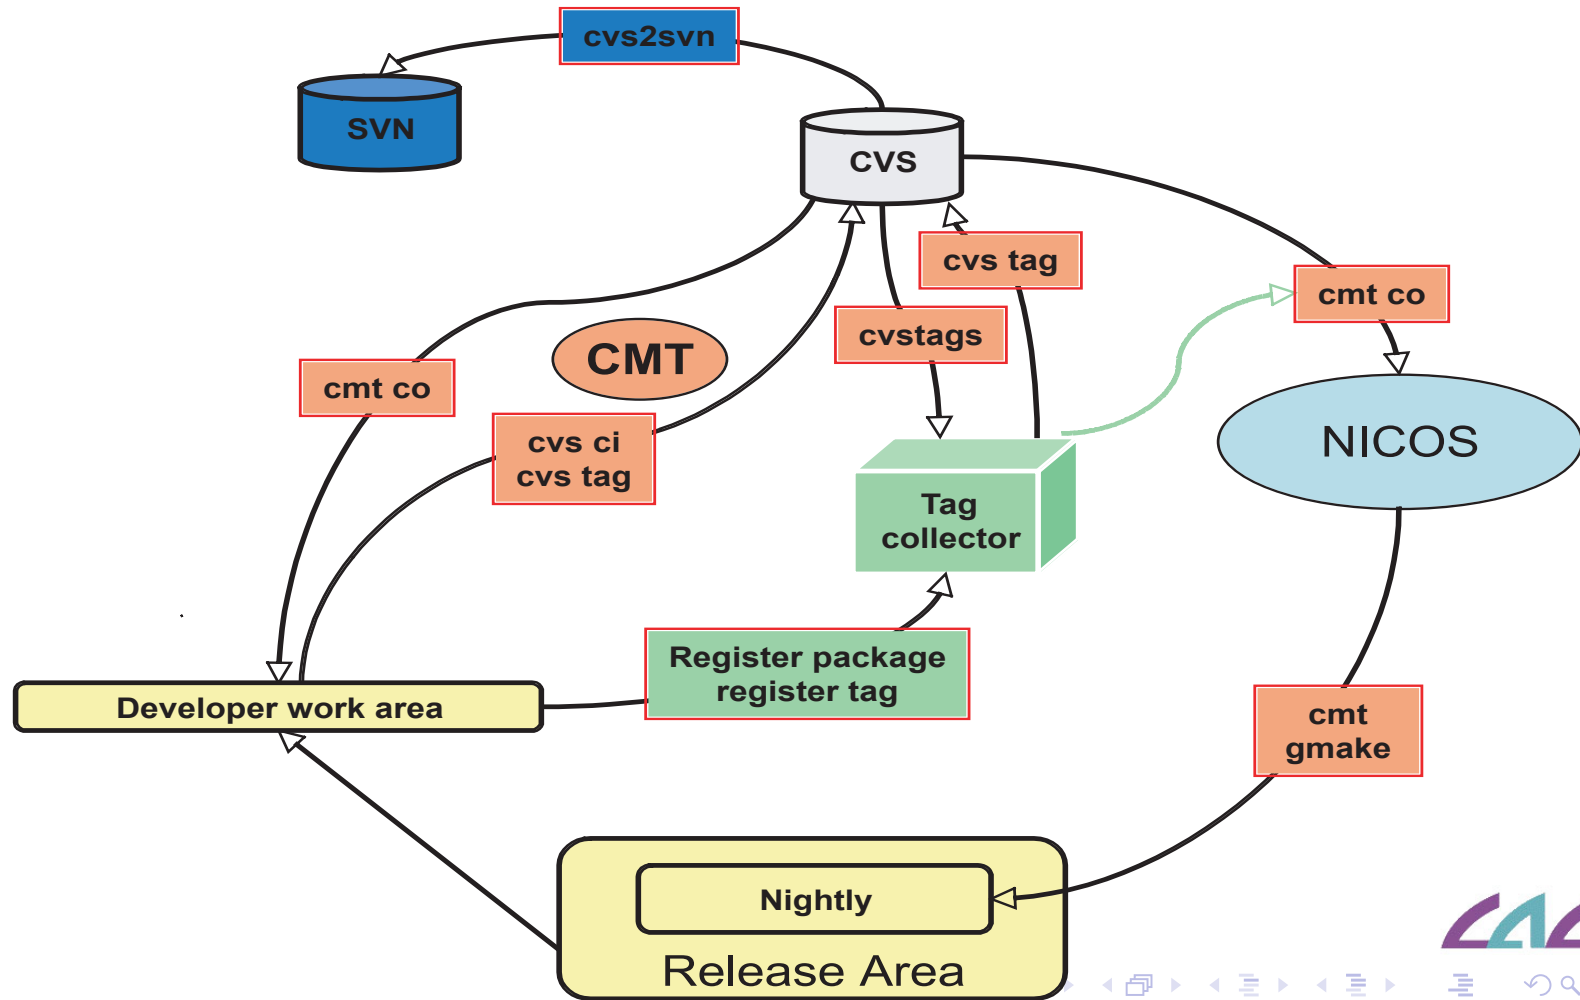


A brief summary by pictures...



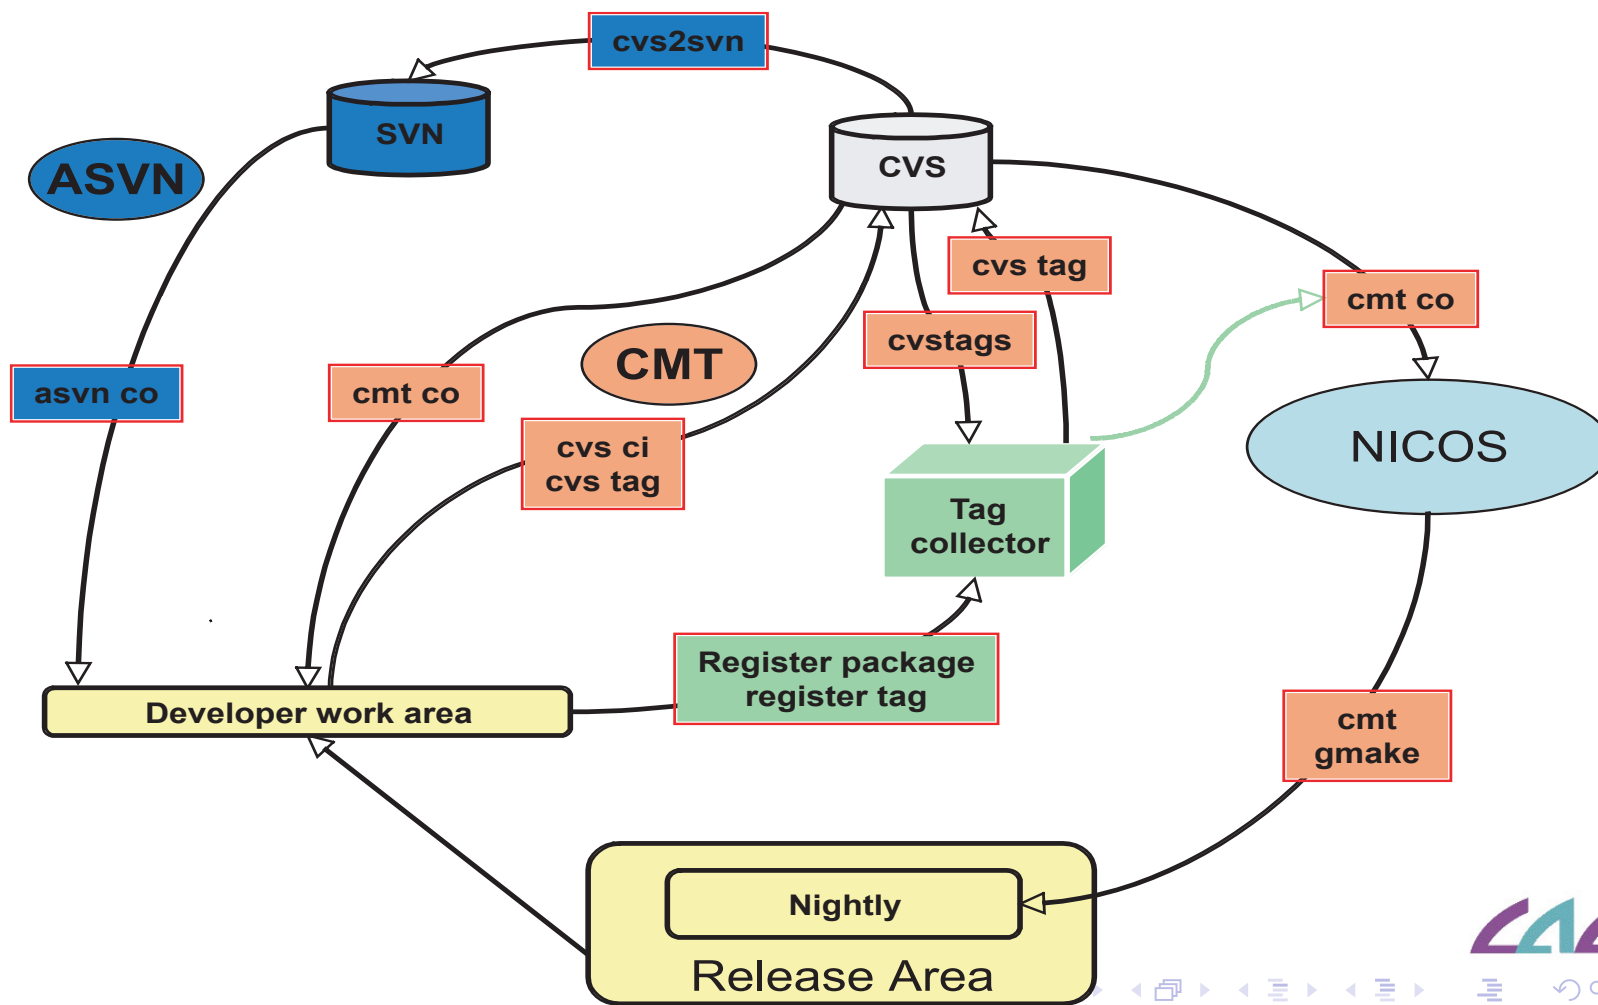


A brief summary by pictures...



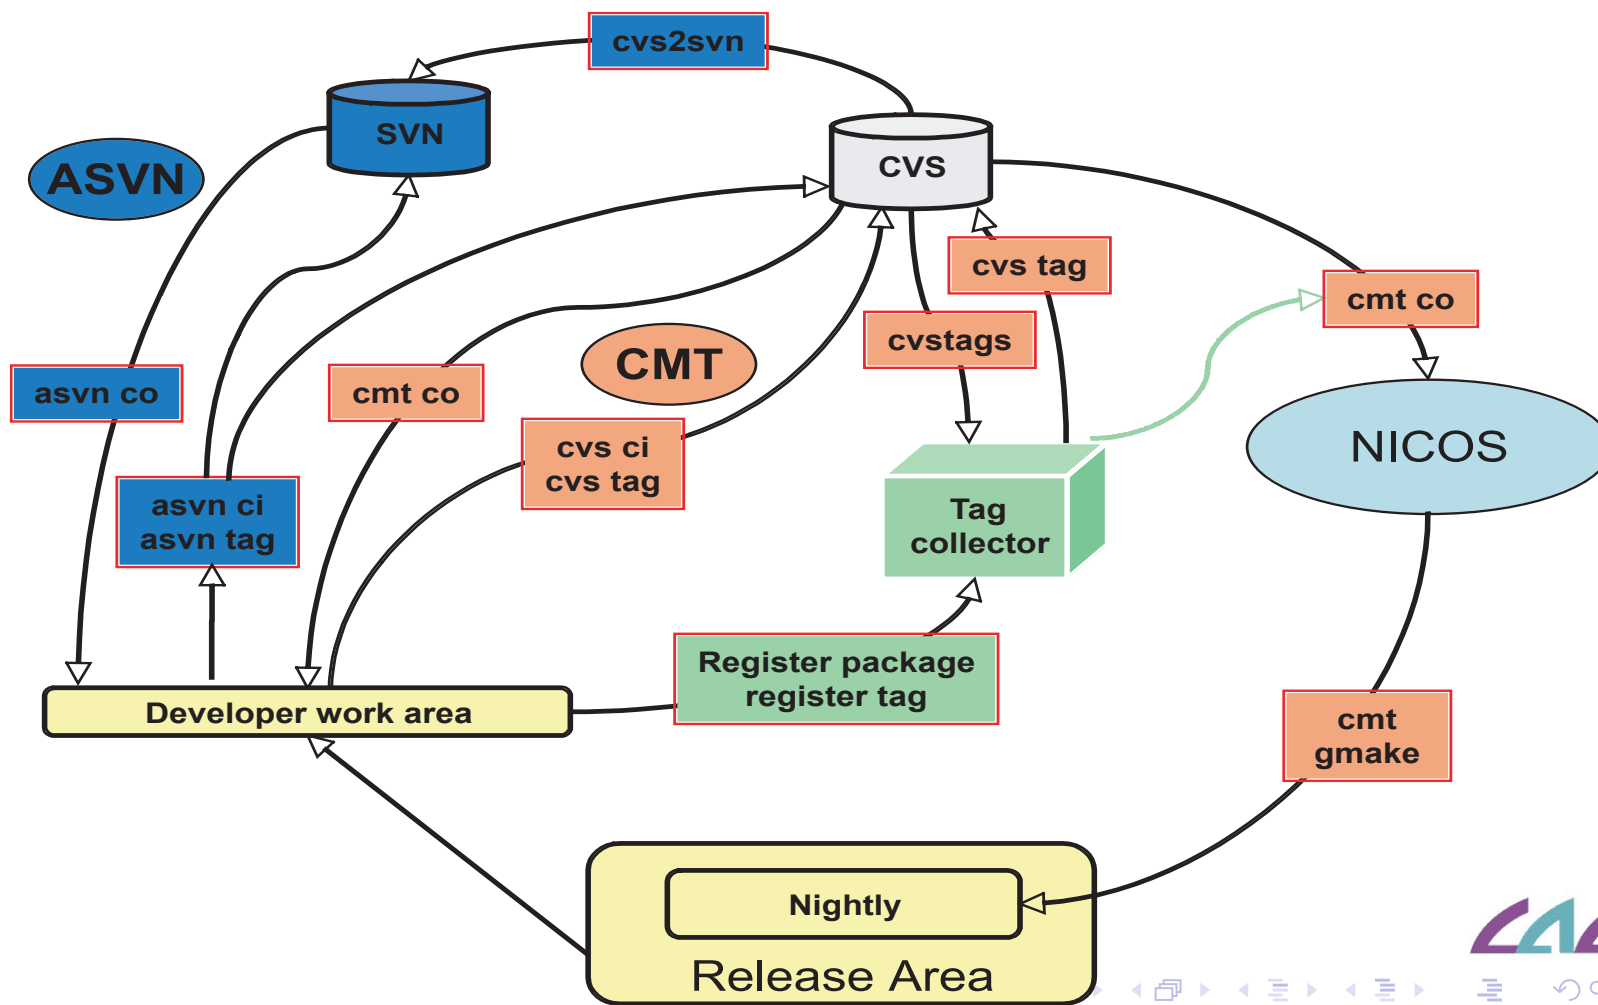


A brief summary by pictures...



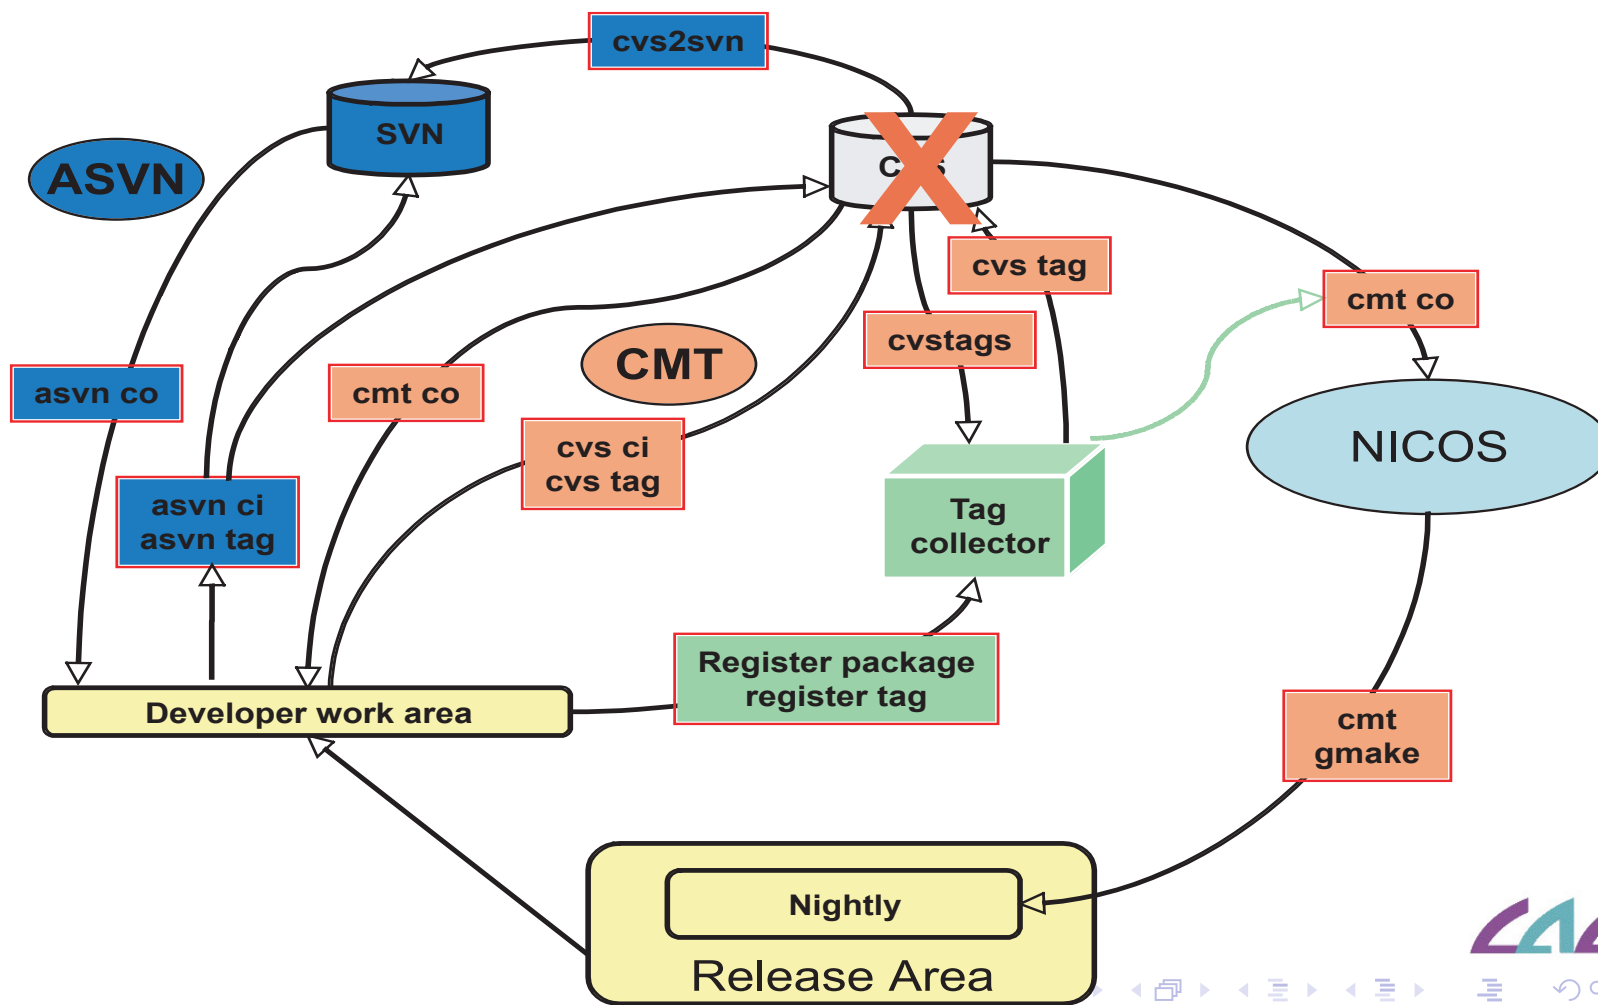


A brief summary by pictures...





A brief summary by pictures...



Future Plans

- I have proposed to ATLAS management that we migrate to SVN following the 2009 Physics Run (i.e. Dec 2009 to Feb 2010)
- This is just a proposal at the moment which needs to be adopted by ATLAS management
 - Not just a software issue since impacts physicist users as well as developers
 - However, initial feedback from Software Management is positive
 - Next stage is to discuss and decide at Computing Management level and finally at Trigger/Computing Oversight Board
- Dust off the earlier proposal and see if there's anything changed that needs it to be modified
 - Don't intend a total re-evaluation
- Note ATLAS already adopting SVN for document management