Outcome of the CMS survey to WLCG site admins

Many thanks to all the people who filled the questionnaire!

Daniele Bonacorsi (University of Bologna)

Disclaimer

This survey has a limited statistical significance

- Just a call for volunteers, not aiming to systematically cover all sites
  - not all Tier levels are equally represented, also
- Actually intended for **people** not for **sites**
  - multiplicity of answers per site can be 0, 1 or N
- I am considering also answers which are part of incomplete surveys

It gives an interesting insight on opinions, nevertheless

- How site people perceive the work CMS does on their sites
- How site people perceive the work they do at the sites to support CMS

No “golden” conclusions intended to be drawn

- Trying to identify unsatisfactory areas, not-so-obvious mistakes we do, etc
- Be quantitative

The survey

10 questions

- a general part
- a technical part
- a communication part

Expected completion time: 15 minutes

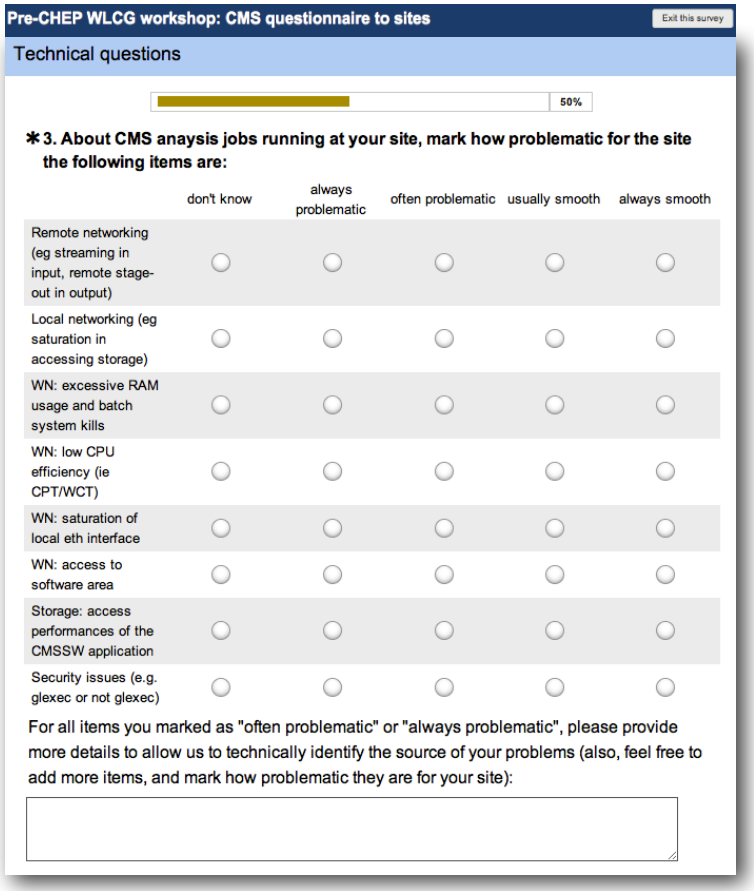

Pre-CHEP WLCG workshop: CMS questionnaire to sites

Exit this survey

Technical questions

50%

✱3. About CMS anaysis jobs running at your site, mark how problematic for the site the following items are:

	don't know	always problematic	often problematic	usually smooth	always smooth
Remote networking (eg streaming in input, remote stage-out in output)	○	○	○	○	○
Local networking (eg saturation in accessing storage)	○	○	○	○	○
WN: excessive RAM usage and batch system kills	○	○	○	○	○
WN: low CPU efficiency (ie CPT/WCT)	○	○	○	○	○
WN: saturation of local eth interface	○	○	○	○	○
WN: access to software area	○	○	○	○	○
Storage: access performances of the CMSSW application	○	○	○	○	○
Security issues (e.g. glexec or not glexec)	○	○	○	○	○

For all items you marked as "often problematic" or "always problematic", please provide more details to allow us to technically identify the source of your problems (also, feel free to add more items, and mark how problematic they are for your site):

Number of initiated surveys: **100**

Number of completed surveys: **41**

- in the next slides, the size of the response pool is shown for each plot
  - look for this: Answers: X

General information

What Tier level?

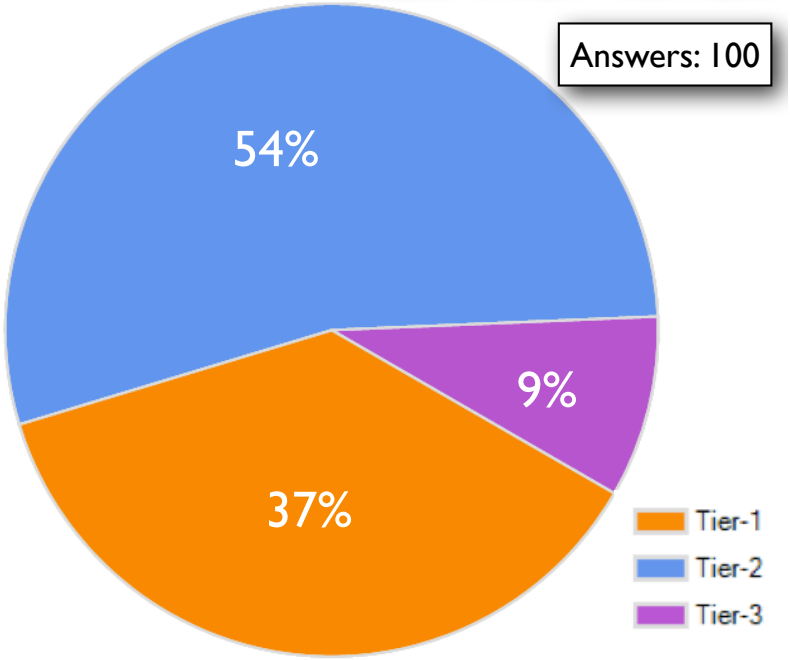

More than 50% of answers come from T2s

- more than expected?
- a good response by T1s, but that was expected
- T3s seem to knock at the WLCG door...

Where?

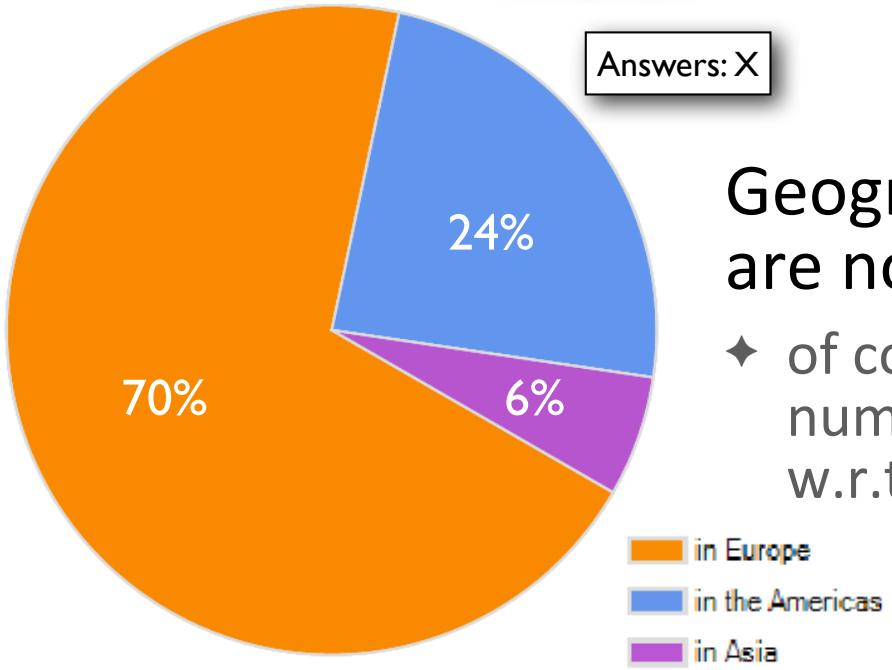

Geographically, individuals who answered are not equally distributed

- of course, it may also be just due to the lower numbers of CMS T2s in the Americas and in Asia w.r.t Europe

About CMS **analysis jobs** running at your site, mark how problematic for the site the following items are:

Answers: 48

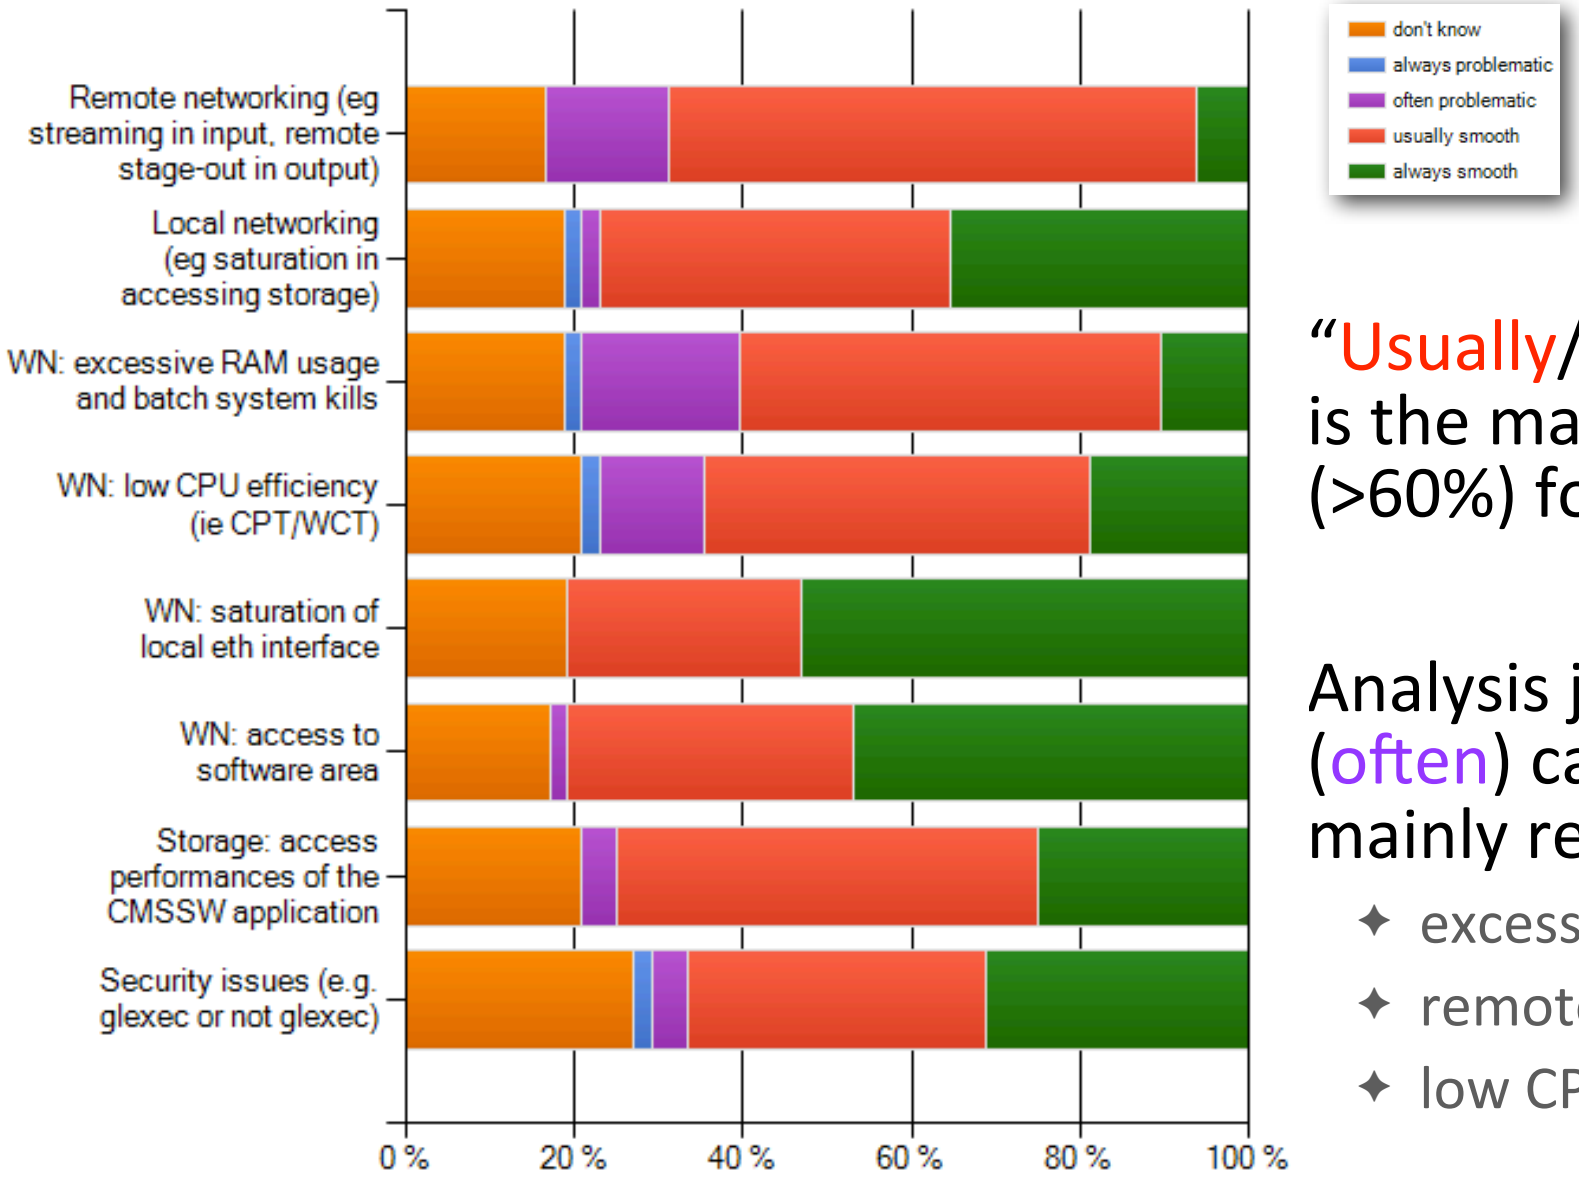

"Usually/always smooth" is the majority answer (>60%) for each item

Analysis jobs seem to (often) cause problems mainly related to:

- excessive RAM usage
- remote networking
- low CPU efficiency

About CMS **Monte Carlo production** or **Reprocessing/ Skimming jobs** running at your site, mark how problematic for the site the following items are:

Answers: 48

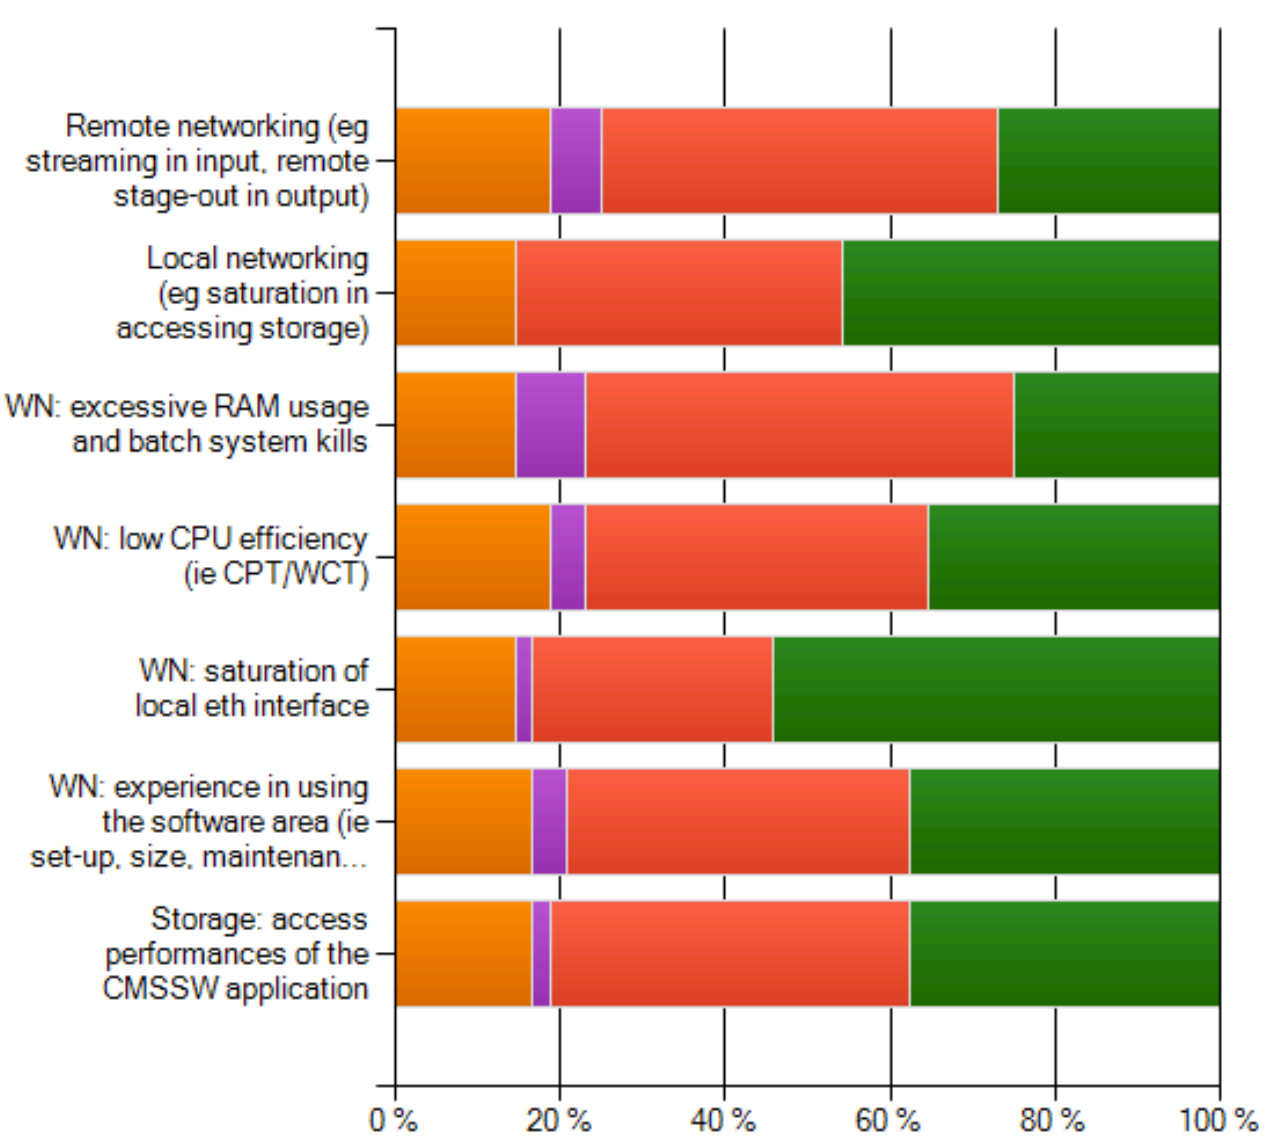

"Usually/always smooth" is ~80% of the answers for each item

- Scheduled processing more stable
- Remarkably, the "always problematic" *never* shows up

These jobs (only "often", never "always") can cause issues

- and same sources as before, RAM usage larger

About CMS **Transfers**, mark how problematic for the site the following items are:

Answers: 48

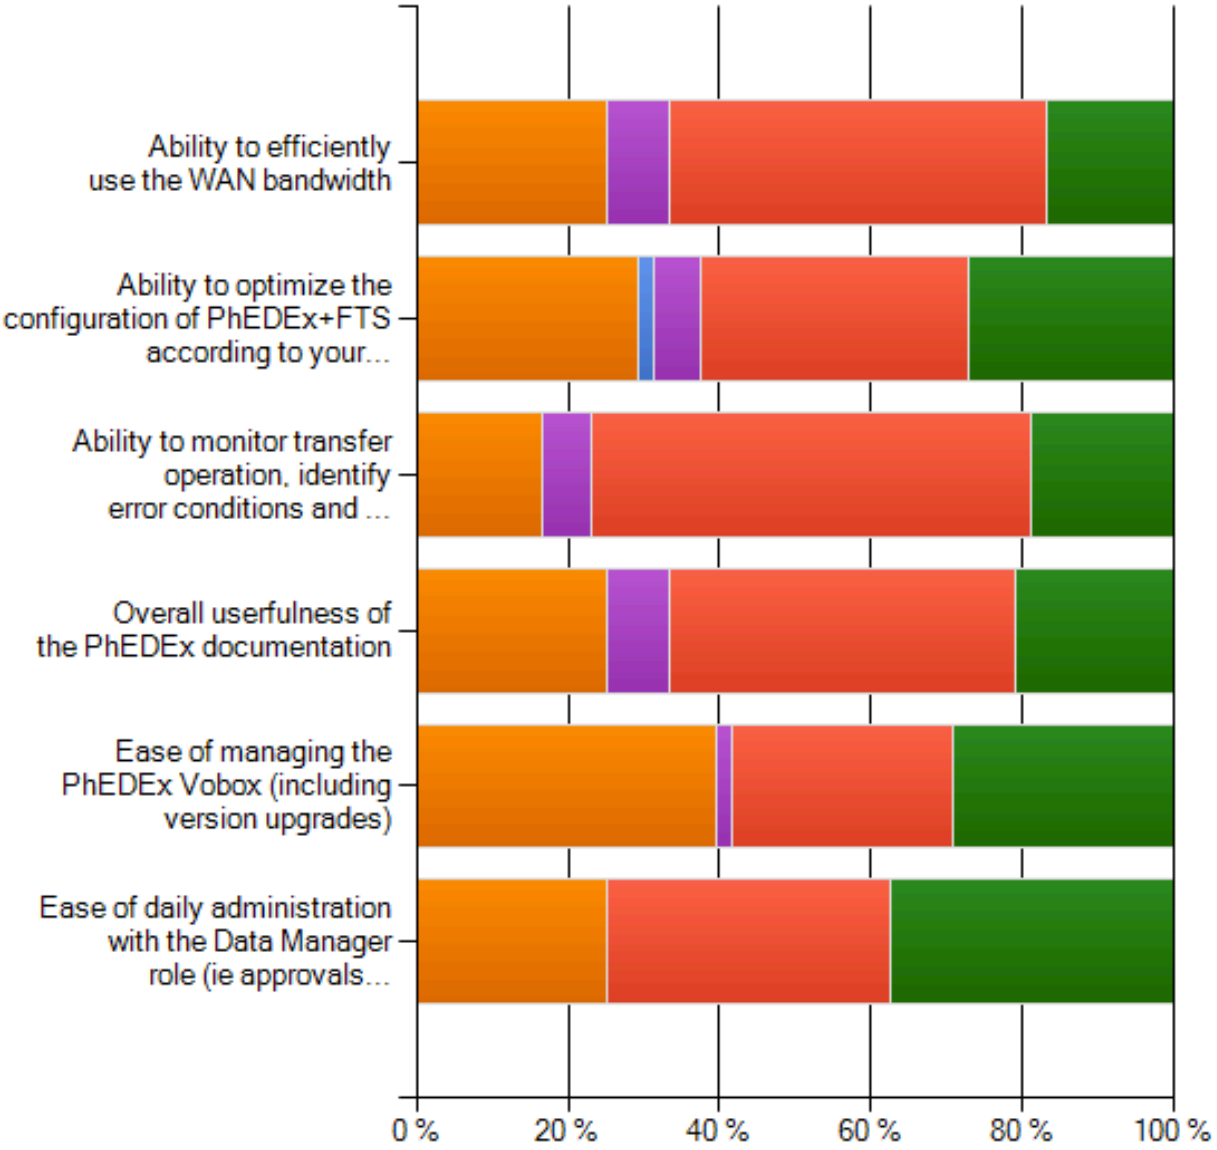

"Usually/always smooth" at least >60% on each item

- remarkably, the "always problematic" rarely shows up for Transfers

Perhaps an action can be triggered, here:

- ease the PhEDEx Ops for sites
  - could the key just be to improve the existing documentation and adapt it for quick and efficient Ops?
- A parameter here should also be that new sites/people have a learning curve

CMS and site monitoring: Availability of efficient tools for *PASSIVE* monitoring of the following:

Answers: 48

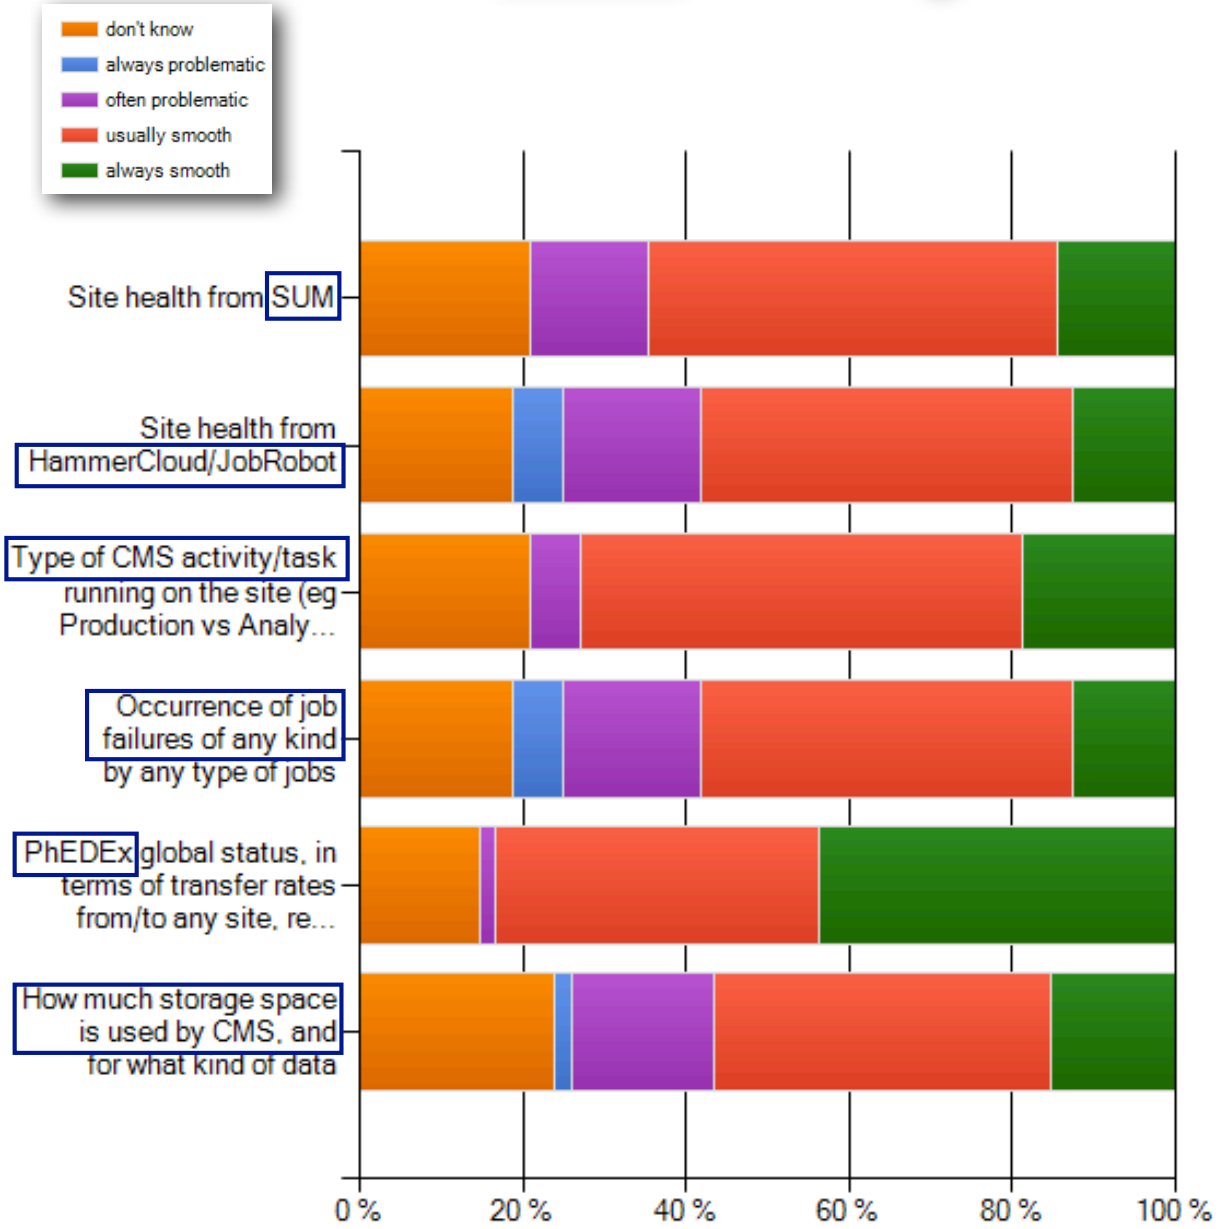

"Usually/always smooth" still the majority, but some structure is visible...

In particular, with the current tools, it seems often/always harder to monitor:

- site health with HC/JR
- occurrence of job failures
- amount of storage used

and often harder to monitor

- site health with SUM

On the other hand, PhEDEx monitoring and activity monitoring (e.g Prod vs Analysis) are quite smooth

Is this unbalance in the relative "usefulness" of the tools available perceived (aka lack of knowledge on how to use the tool) or real (aka tool inadequate to meet the users needs)?

CMS and site monitoring: Availability of efficient tools/support for *ACTIVE* monitoring of the following:

("*ACTIVE*" == these data will be used to take actions, not just to check the general site health)

Answers: 48

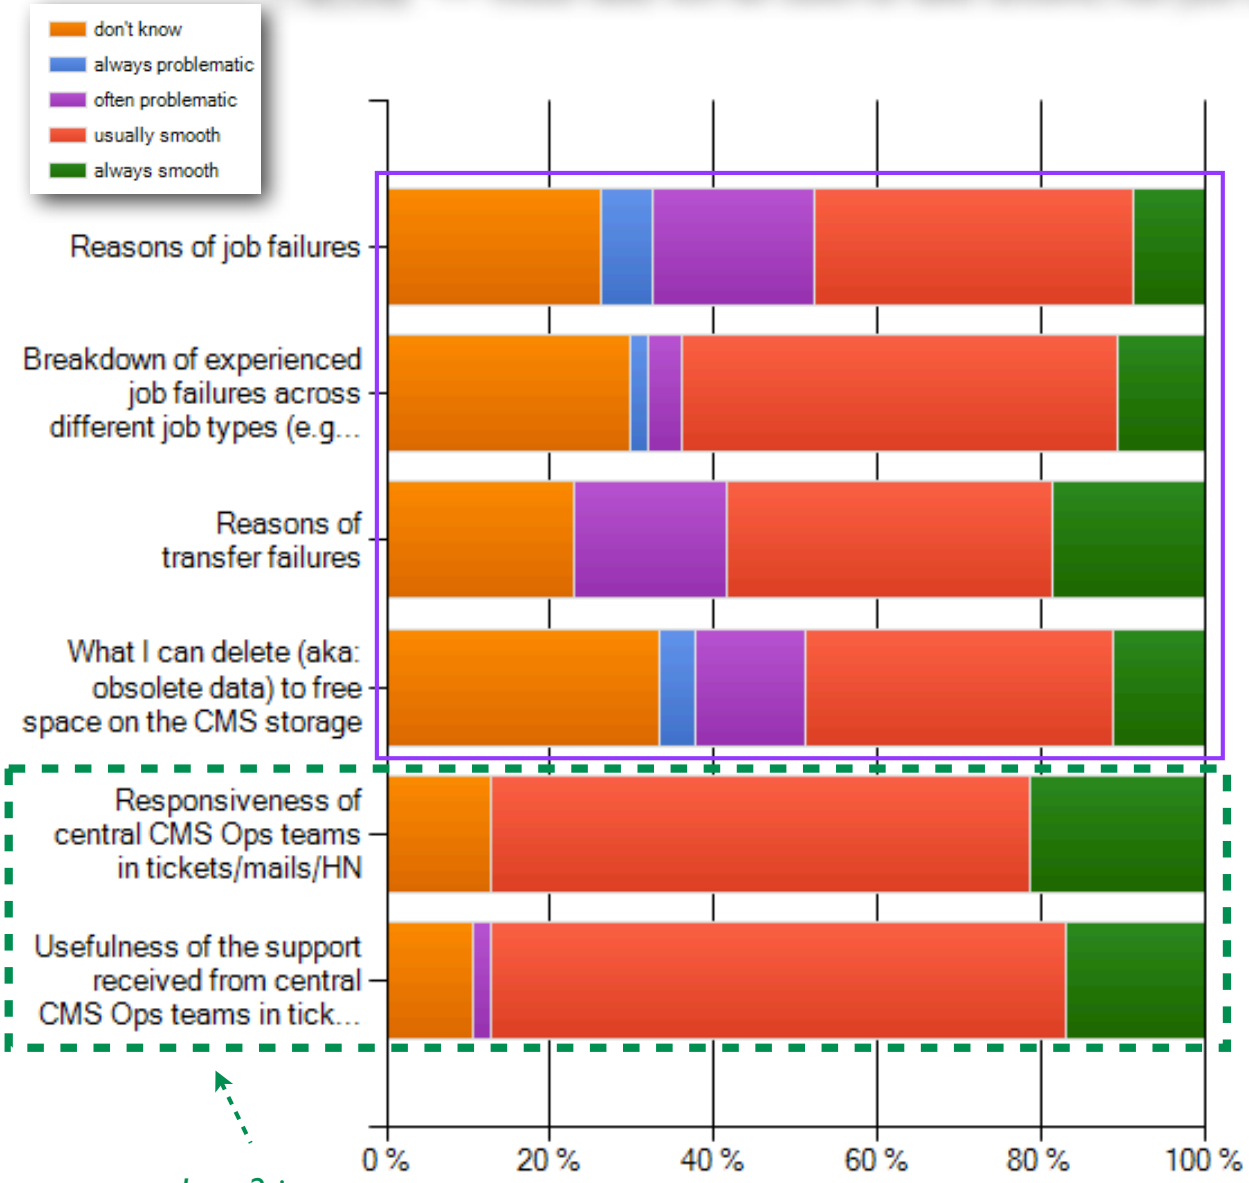

Last 2 items: slightly different questions

Quite some structure here..

Finding the reasons of job/transfer failures is "problematic" in 20-25% of answers

- interesting for the Transfers: the PASSIVE part was ok...
- my guess: it's because the sources of effectively useful information are diverse, and there is no unique and unambiguous access to such info?

We can improve in tagging data as obsolete

- my guess: more an organizational/ communication issue than a technical one?

Nevertheless, the responsiveness of the CMS Ops teams is smooth and the support received is useful (~90%)

My (rough) summary: a step forward in the "active" monitoring could still potentially help the sites to help themselves more: the CMS support teams anyway work well with the sites and allow Ops to quickly recover

General level of satisfaction about **how CMS and the sites communicate among each other**, in terms of:

Answers: 46

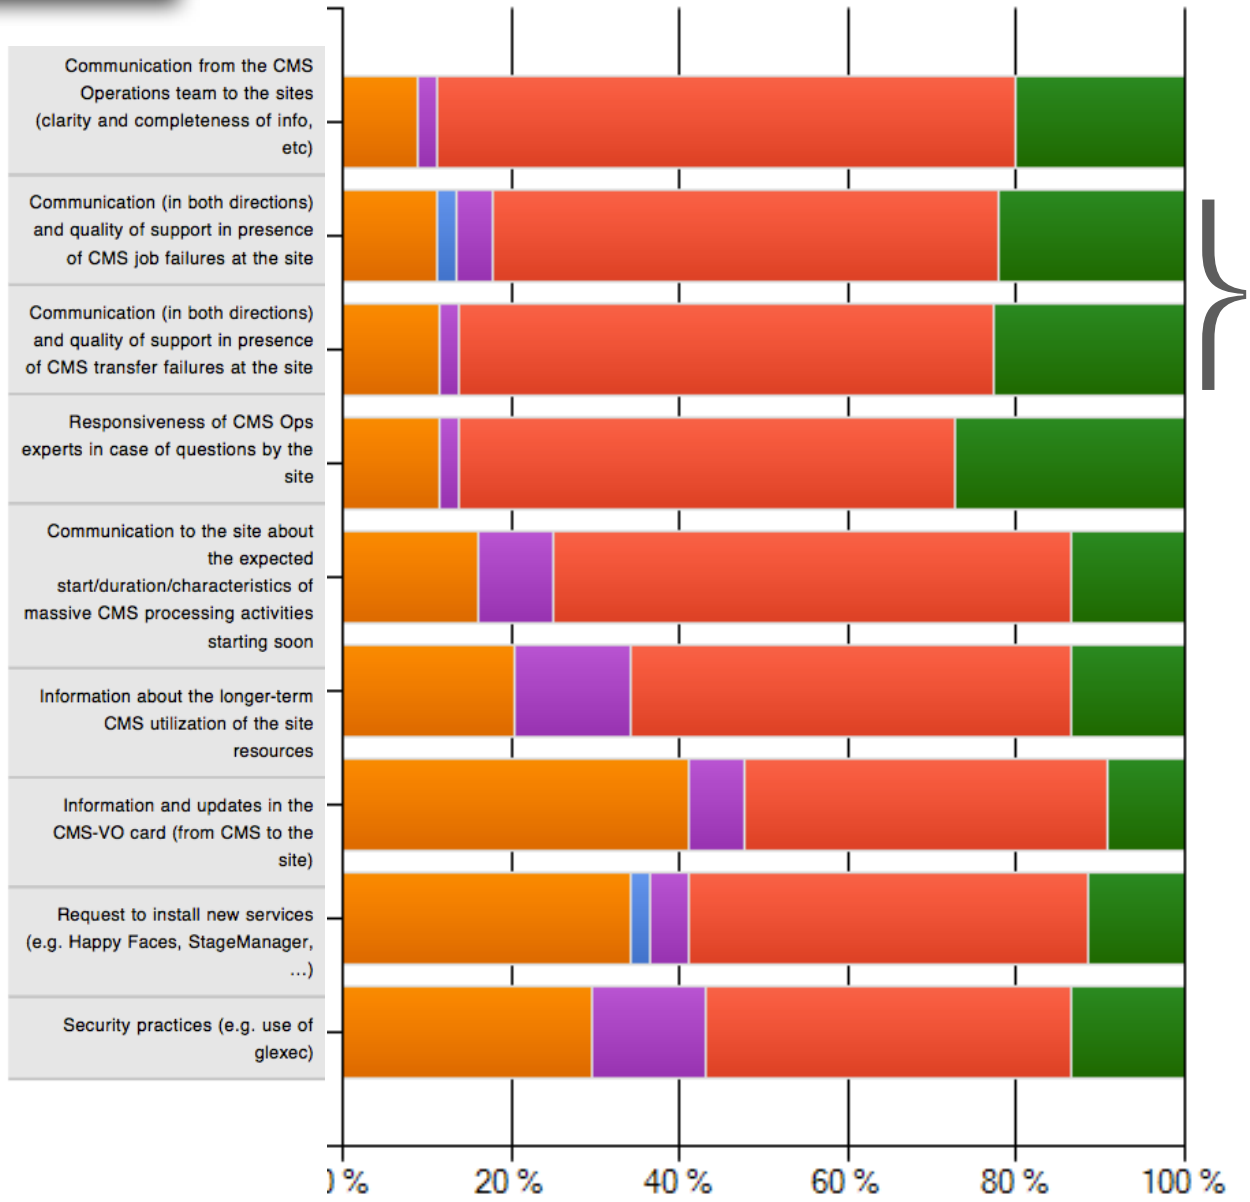

Site are satisfied about how CMS Ops talks to sites at the 90% level

Comparably good (>80%), with a slightly better support obtainable for transfer issues

OK

Computing Ops should communicate more specific and detailed info to sites? More direct T2 representation at CMS Ops meetings?

Same as above, even more visible... Some frustration maybe. Margin of improvement to be explored

40% of "don't know" in the CMS VO card!?

I would have expected sites to be more against this. Guess: probably it really depends on which service, and how they manage to learn how it works beforehand?

Generally not bad, maybe some frustration though?

Answers: 41

In general, **my experience** in collaborating with CMS at my site during Run1 so far was:

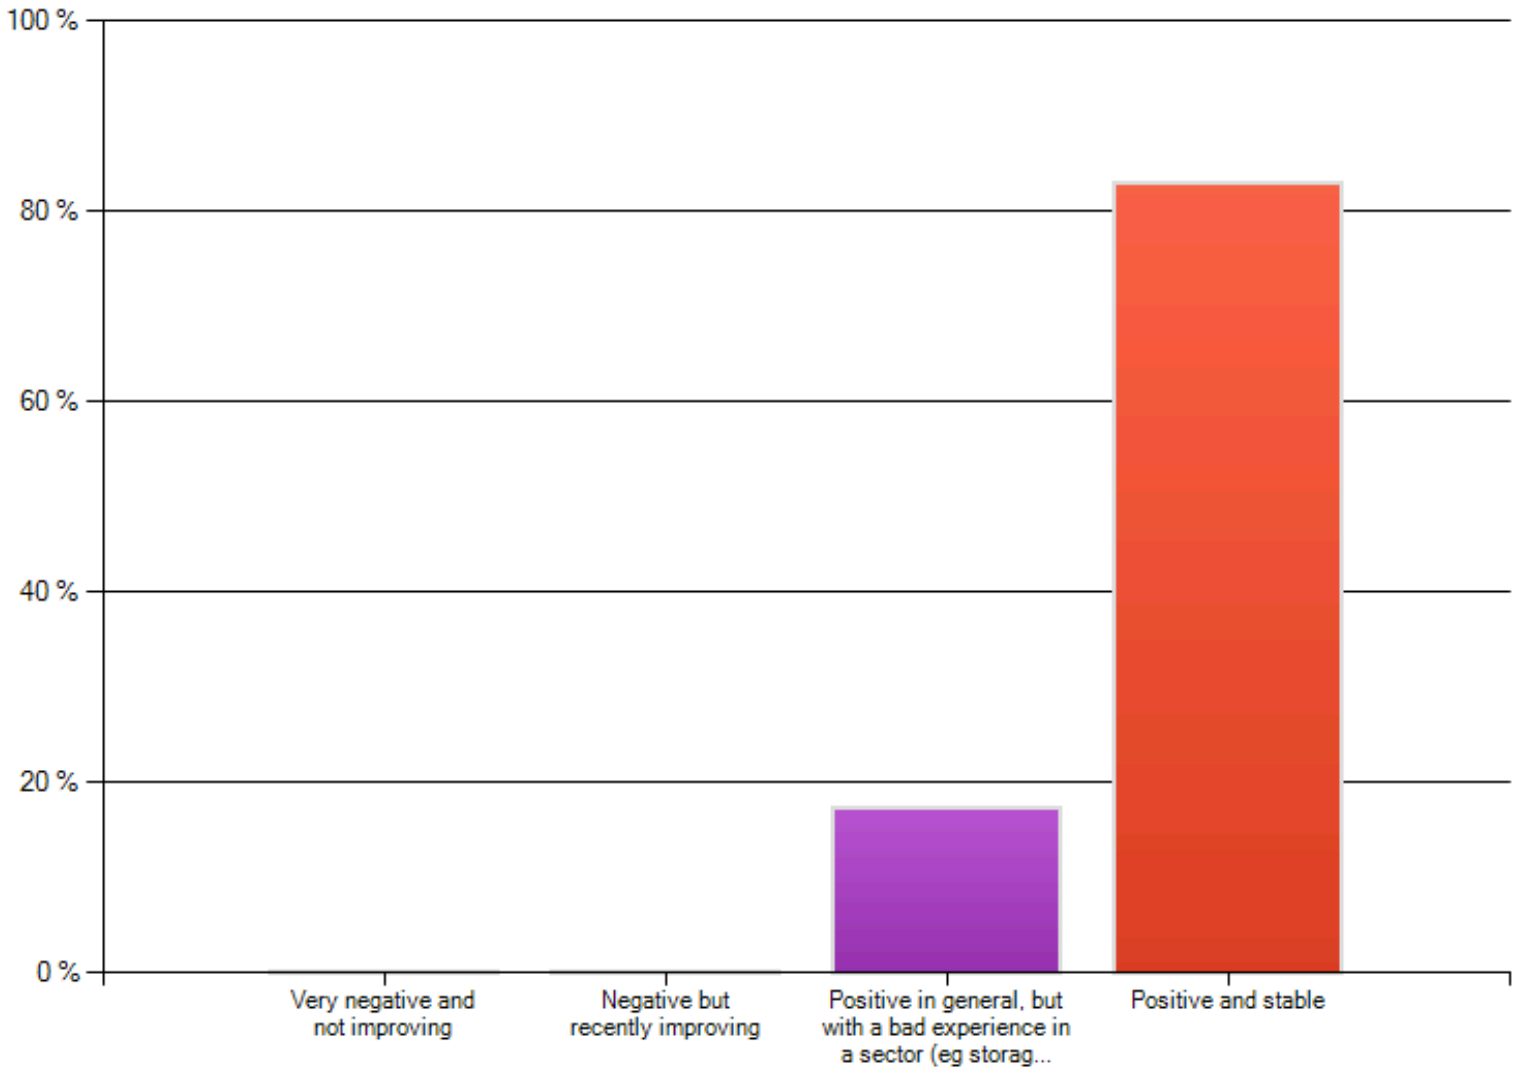

Other items, problems, suggestions you want to share with us (open text)

Answers: 41

(positive comments are intentionally omitted)

1. SUM monitor pages. All the recent changes have introduceed more dynamical pages, with the use of ajax and similar technologies, resulting in pages more estetically pleasing but less practical. In our opinion a simple plain HTML would be more useful. 2. Exploitation of grid information system and related databases (SIteDB, GOCDB,..). The new tools introduced to operate the MC production don't appear to use the dynamical information available and published by sites. For example, when a site adds or removes a CE it's now necessary to open a ticket and have someone to manually update some list. Likewise, the load of jobs on multiple CEs is not uniformly spread among them.

Highlighting 3 comments by different individuals that show how important the feedback from sites is:
we should listen more!

Working very closely with CMS our site is able to provide the best service. If there is an area of improvement it is in (1) occasional analysis jobs wanting to use unreasonable amount of memory (2) "stale" or "rather unused" data sets stored on our site taking up large chunks of space and (3) some times file transfer connections between various sites requires special action for upkeep.

Information about disk hardware that works well would be interesting to share among the sites. CMSSW should be ported faster to new OS versions to enable sites to use more recent hardware efficently. Documentation is important to keep up to date and should be easy to find. Presently some site related pages are "lost" in the TWiki and there is no simple way to access all pages from one place.

Nothing new: documentation is always lagging behind and in any case needs a robust consolidation We badly need a unique, up to date, tree of docs rather than an hypergraph of twiki pages

Some other interesting "trends" in comments

"job monitoring is harder than transfer monitoring"

Would be good to have tools for more detailed analysis of failures in jobs and transfers.

So far so good :-) Just to re-state that, while the data transfers system is very good for sites to find information about performance and issues (phedex), the same can not be said about job management. The fact that several systems are (...)

WMAgent-based jobs are difficult to debug from a site perspective. The stdout/err/logs available to the CompOps folks are rarely sufficient to figure out what's going wrong.

The ATLAS web interface for looking at job details, in particular for accessing information quickly about job failures (worker node, CE, error messages, ...) seems far superior to anything CMS has.

"Lack of control in outbound traffic"

Maybe, In the nowadays network configuration the FTS is a "bottleneck" for several T2-T2 transfers. Using the Phedex agents, you have control about the incoming jobs/transfers, but you cannot control the outcoming transfers running on your site, and a overload on the gridftp service is more usual than recommended

"GGUS:Savannah = 1:0"

Please use GGUS instead of Savannah (bridge to GGUS does not work everytime)

"Explore ways to better communicate the plans"

Receive more details on software under testing and timescales