

Automating usability of ATLAS Distributed Computing resources

Salvatore A. Tupputi (INFN CNAF) on behalf of ATLAS Collaboration



Distributed Processing and Data Handling A:
Infrastructure, Sites and Virtualization
14/10/2013

Outline

- Automation in ATLAS Distributed Computing (**ADC**) operations
- Storage Area Automatic Blacklisting (**SAAB**) principles and implementation
- Experience with Site Availability Monitoring (**SAM**) tests for storage resources
- SAAB in **production**: monitoring and automation **experience**
- Further ahead: short and long term **plans**, then **conclusions**

ADC and ADC automation overview

Site Status Board (SSB)

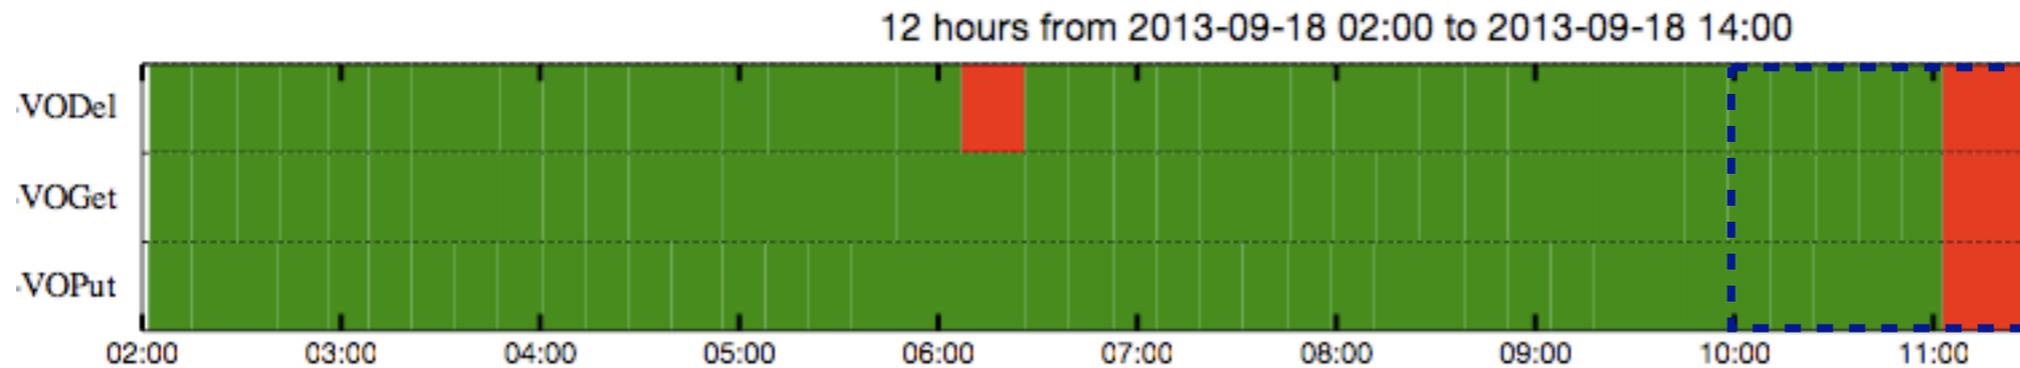
- For all the ATLAS computing sites ADC attends to **manage**
 - **resources**: storage and computing elements, network, catalog, File Transfer System, ...
 - **activities**: distributed data processing and analysis, data transfers
- Tools for **automatic** control are being introduced in several ADC areas:
 - *Switcher* for automatically testing and recovering queues after downtimes
 - *Panda site exclusion* for automatically handling sites computing resources
 - Automatic exclusion/recovery of storage resources for *downtimes* or *disk space* quotas reaching critical thresholds
- Several tools provide **monitoring** information which is uniformly collected in the SSB
 - SSB views and SSB metric allow to follow sites status and history according to given metrics

Distributed Data Management (DDM)

- ADC storage resources are accounted for by **Distributed Data Management (DDM)** system
- The DDM End Point (**DDM EP**) represents the elemental ATLAS **storage unit**
 - Each ATLAS site has one or more of them
 - DDM system keeps track of the status of DDM EPs:
 - each of them can be in **on/off/auto** mode
 - only '**auto**' DDM EPs can be automatically acted upon
- Automatic management of DDM units take into account
 - Storage downtimes, space shortage
 - Performance-related metrics missing: that's where SAAB's idea has started

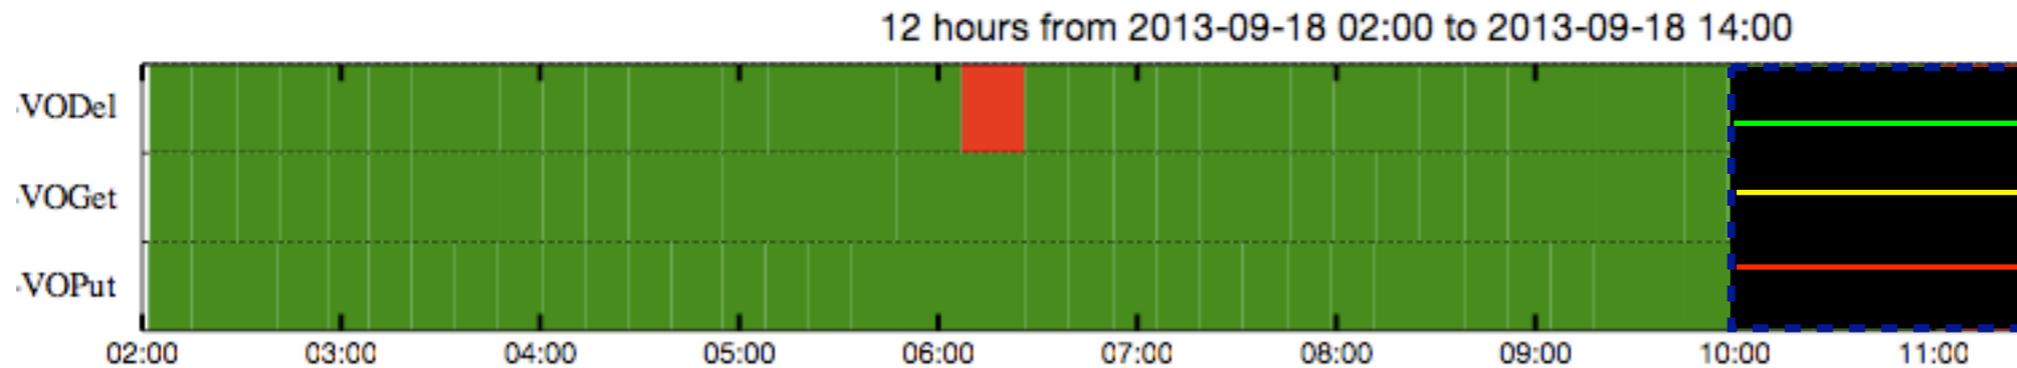
SAAB concepts

- Given a metric regularly yielding results in time, define a **fixed-width** time **interval** taking into account several of them



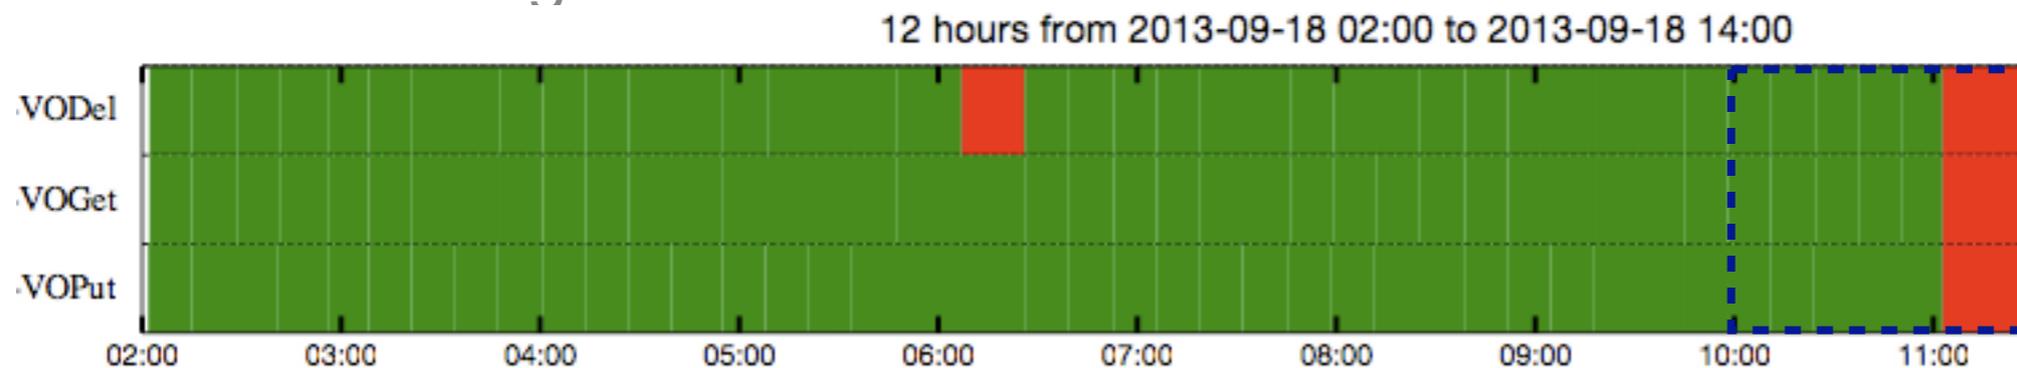
SAAB concepts

- Given a metric regularly yielding results in time, define a **fixed-width time interval** taking into account several of them
- Define and deploy an **algorithm** providing performance-defining **thresholds**



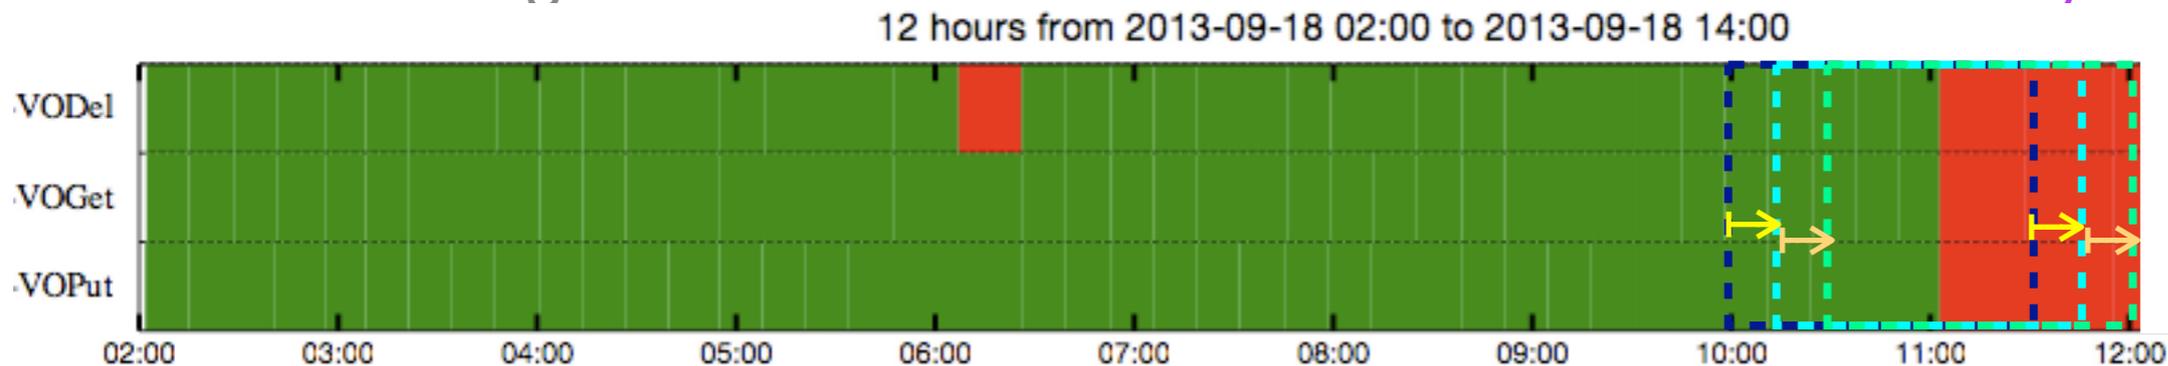
SAAB concepts

- Given a metric regularly yielding results in time, define a **fixed-width time interval** taking into account several of them
- Define and deploy an **algorithm** providing performance-defining **thresholds**
 - according to which a **decision** is issued to be **automatically** taken



SAAB concepts

- Given a metric regularly yielding results in time, define a **fixed-width** time **interval** taking into account several of them
- Define and deploy an **algorithm** providing performance-defining **thresholds**
 - according to which a **decision** is issued to be **automatically** taken

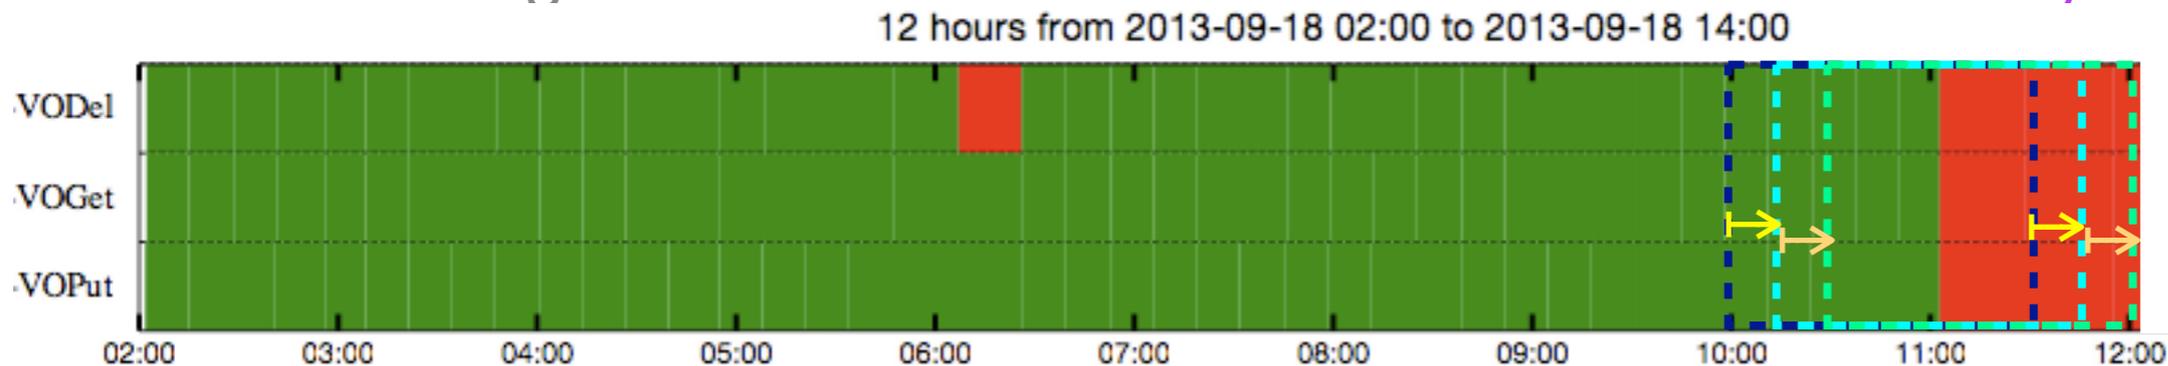


- **Repeat**, by shifting the time window by constant time-steps



SAAB concepts

- Given a metric regularly yielding results in time, define a **fixed-width time interval** taking into account several of them
- Define and deploy an **algorithm** providing performance-defining **thresholds**
 - according to which a **decision** is issued to be **automatically** taken



- **Repeat**, by shifting the time window by constant time-steps
- This approach provides a **sliding time-window** over the metric to be looked at
 - Ground is laid for **automatic blacklisting** of resources whose behavior is found not compliant with a suited blacklisting policy

SAM tests probing sites resources

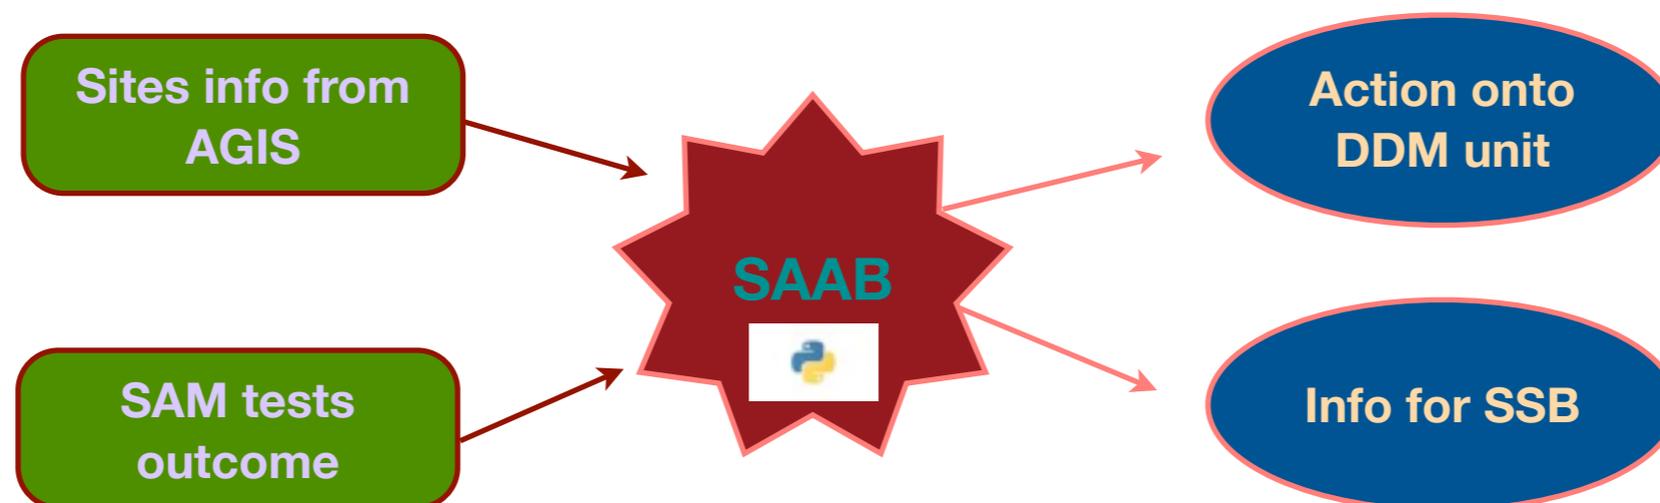
- SAM is the **framework** used by the experiments to **test** status and availability of resources
 - It is a **Nagios**-based service, tests have 3 possible **outcomes**: **OK**, **WARNING**, **CRITICAL**



- The most relevant tests for storage resources management are **Put**, **Get** and **Del** of a test file into the storage area (DDM EPs)
- SAM framework probes ATLAS sites **availability** and the results are stored into a database
 - **latest** results and **historical** data for each site are exposed through a **RESTful** interface
 - Tests are executed every **13 minutes**
 - This is the highest frequency allowed by the used machines HW

SAAB applied upon SAM tests in ATLAS

- SAAB tool first application
 - **processes information** of SAM Put tests and **issues** a set of **decisions** for all the DDM EPs reflecting their **status**
 - **blacklisting/whitelisting** actions for write/upload **permissions** on DDM EPs are taken according to such decisions
 - set to be executed every **15 minutes**, so that each new execution has at least one new SAM test on average as input
 - exposes **output** compliant with the **ATLAS Site Status Board (SSB)**
 - implemented in *python*



SAAB algorithm

- Thresholds and time values have been tuned after a testing experience
 - minimizing fluctuations due to temporary glitches and improving SAAB response stability
- The time window length is set to **90 minutes** (customizable), hence **~6** expected tests per window (**ETW**).
- A **string flag** is assigned to each **DDM EP** (thresholds customizable):
 - **RED** (off) if 0% of the total tests number is OK (**blacklisting** candidate)
 - **YELLOW** (inter-threshold) if the OK tests are below 60% of the total tests number
 - **GREEN** (on) if > 60% of the total tests number is OK (**whitelisting** candidate)
 - **No less** than **30% ETW** is **required** for each site for the algorithm to be applied and action to be taken
 - if not the DDM EP is flagged as GRAY
- String flags are chosen in compliancy with the SSB color convention



SAAB actions with SAM results input

- Every 15 minutes SAAB

- processes, for every DDM EPs, the SAM results within the last 90 minutes
- retrieves information about the already blacklisted DDM EPs

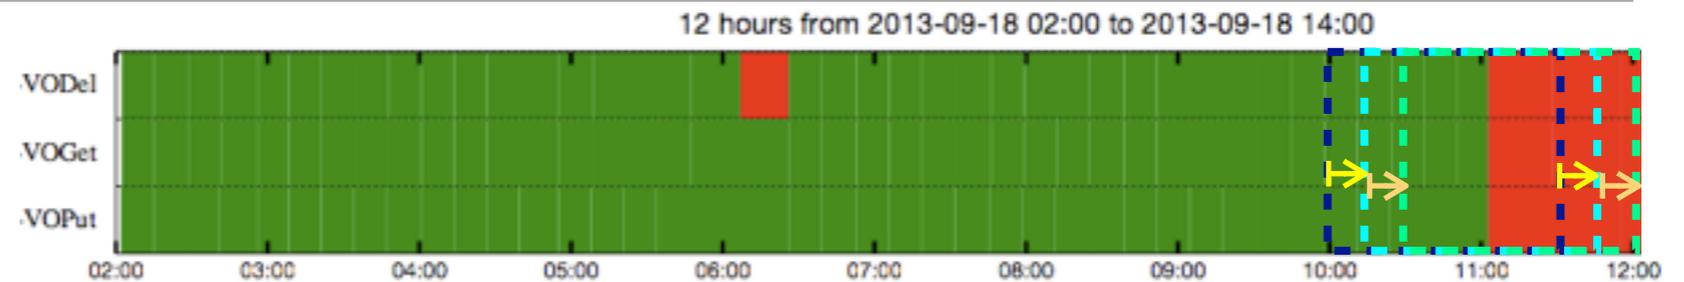
- Blacklisting/whitelisting actions follow via sending instructions to appropriate system in DDM (DQ2):

- RED-flagged DDM EPs are set off for 90 minutes, regardless of their current status
 - If no other change happens, DDM EPs stay off until the 90 minutes expire
 - SAM tests keep being executed anyway
- GREEN-flagged DDM EPs currently off are set on.
- No action taken in any other case (YELLOW, GRAY).

- Information on each execution is written into two logs, accessible via webpage for shifters consumption

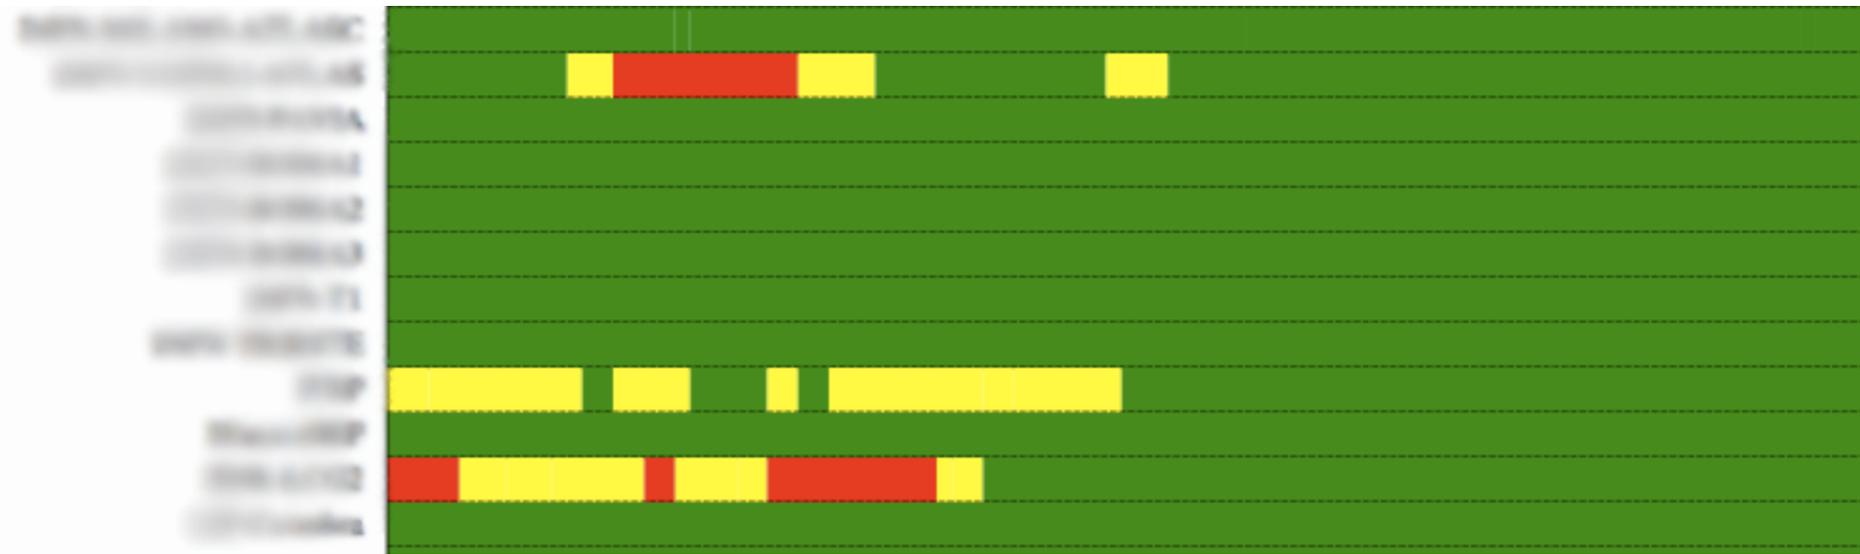
- SAAB sends mail alerts and notifications:

- to developers in case of (known and tracked) problems with SAM results retrieval
- to sites who want to be notified whenever any of their DDM EPs gets blacklisted



SAAB integration in SSB: metrics

- SAAB outcome in **site-grained** fashion showing sites status as the OR of their DDM EPs statuses



- SAAB outcome in **DDM EP-grained** fashion showing DDM EPs status history



- Gray zones relate to SAM tests temporary glitches so that developers can look into that
- Clicking on the site SSB row redirects to the **SAM tests history** for that site (SAAB concepts slide).

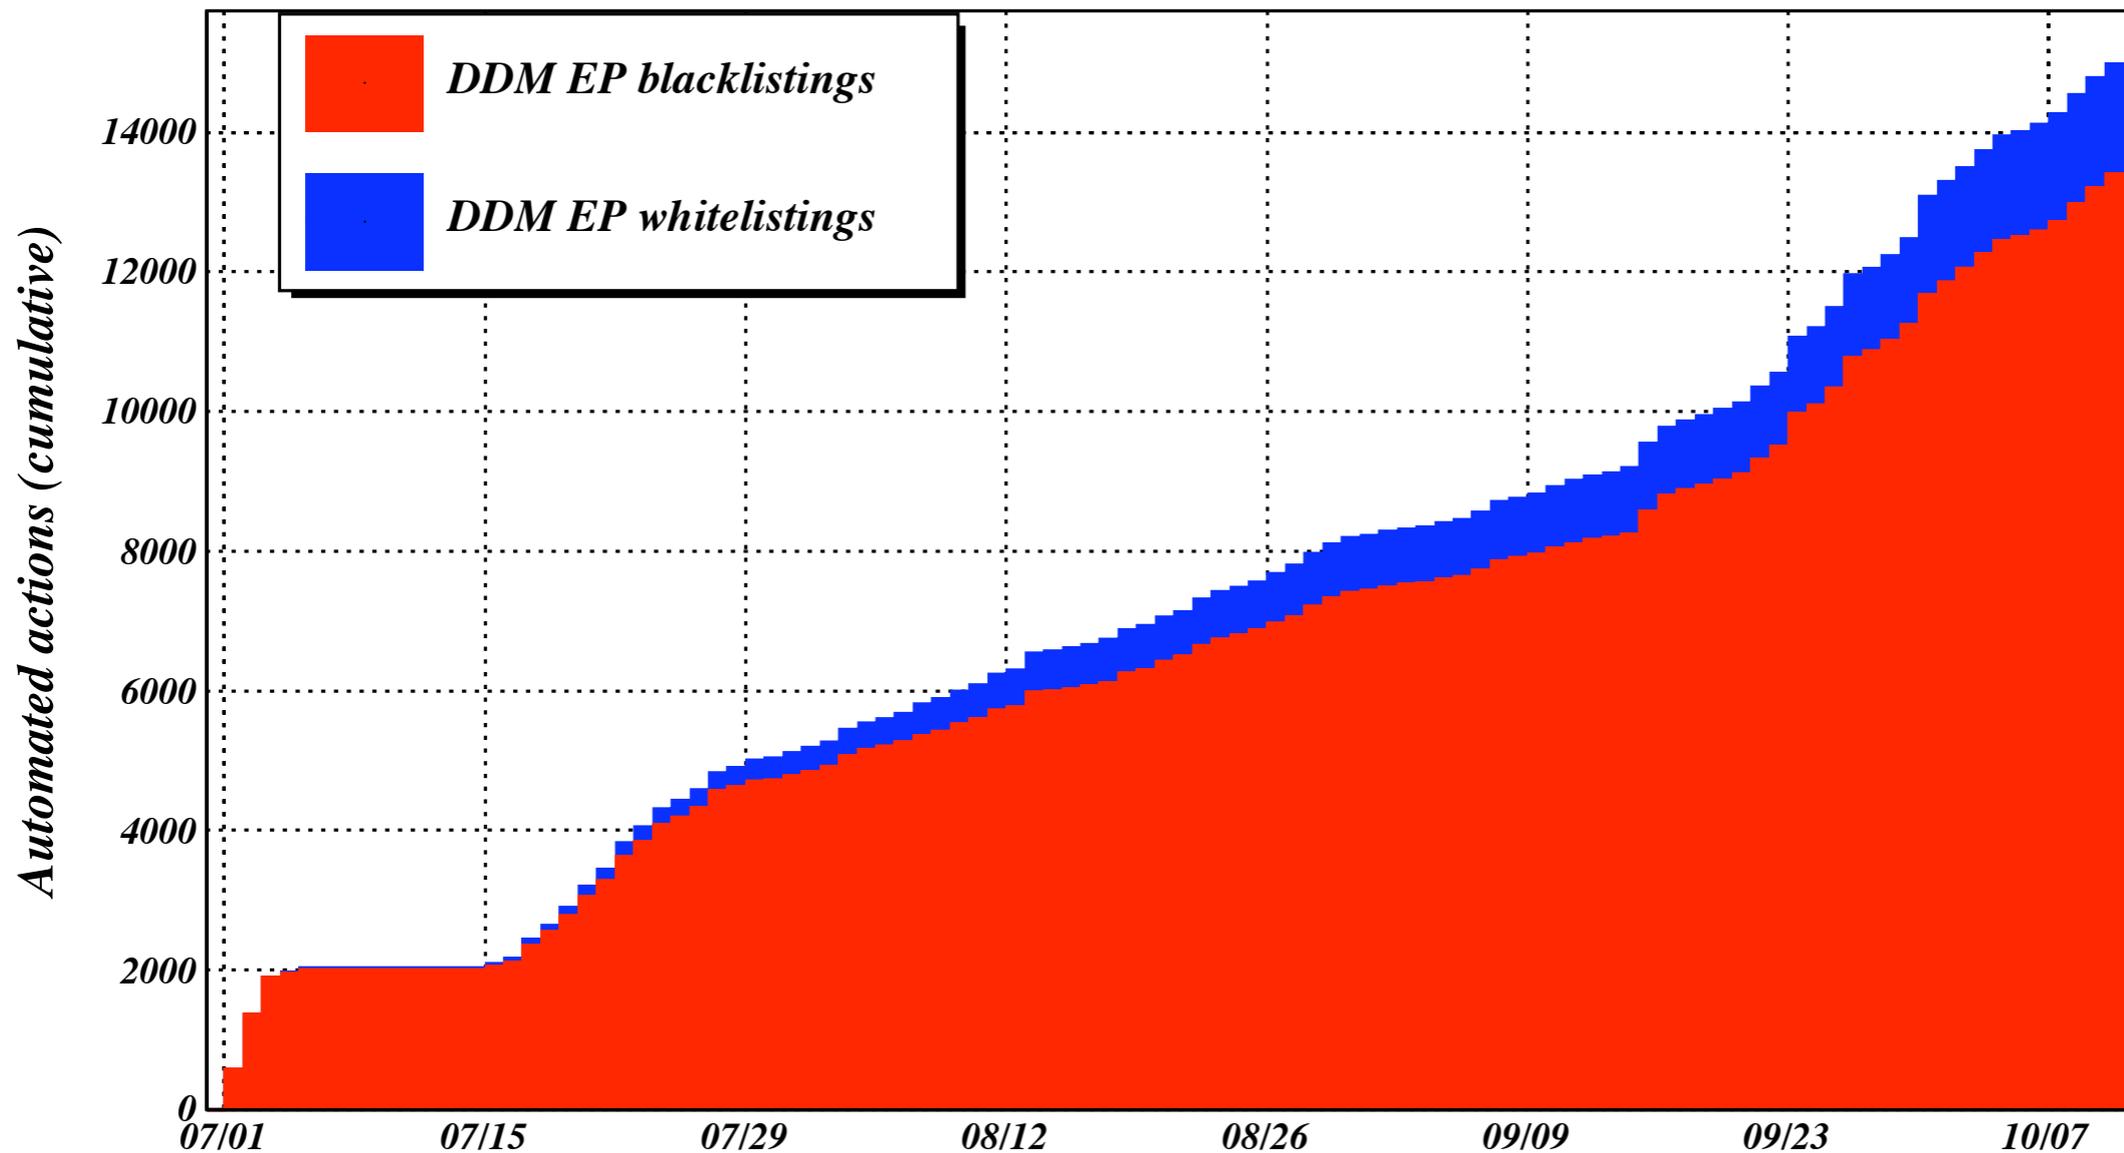
SAAB integration in SSB: views

- SAAB-fed SSB metric is displayed, along with other SSB metrics, in the general SSB **view** for general **shifters**

Site Name ^	Site Info			DataManagement		Functional Tests		Panda Efficiency				Activity status now: Included-Excluded sites				
	Tier ^	Cloud ^	Downtime ^	DDM 4h [%] ^	SITE-OR of DDM upload and write permissions ^	HC_AFT ^	HC_PFT ^	Prod Efficiency 12h [%] ^	Prod Failed Jobs 12h ^	Analy Efficiency 12h [%] ^	Analy Failed Jobs 12h ^	panda prod status NEW ^	panda analy status NEW ^	DDM DA Status ^	DDM DP Status ^	DDM DT Status ^
IN2PS-OPPM	T2D	FR	ACTIVE	100	3/3	100	100	99	4	88	219	online	online	online	online	online
IN2PS-LAPP	T2D	FR	ACTIVE	100	3/3	100	100	99	8	99	7	online	online	online	online	online
IN2PS-LPC	T2D	FR	ACTIVE	100	3/3	100	100	98	10	97	136	online	online	online	online	online
IN2PS-LPSC	T2D	FR	ACTIVE	100	3/3	100	100	89	87	76	88	online	online	online	online	online
INFN-BOLOGNA-T3	T3	IT	ACTIVE	100	2/2	no-test	100	83	51	100	0	online	test	online	online	online
INFN-COSENZA	T3	IT	ACTIVE	100	2/2	no-test	100	98	1	100	0	online	online	online	online	online
INFN-FRASCATI	T2	IT	ACTIVE	97	3/3	0	100	97	14	78	399	online	online	online	online	online
INFN-GENOVA	T3	IT	ACTIVE	100	2/2	100	no-test	88	1	96	14	online	online	online	online	online
INFN-LECCE	T3	IT	ACTIVE	100	2/2	no-test	no-test	100	0	no.data	no.data	online	online	online	online	online
INFN-ILANO-ATLAS	T2D	IT	ACTIVE	100	4/4	100	100	97	21	100	2	online	online	online	online	online
INFN-NAPOLI-ATLAS	T2D	IT	ACTIVE	100	3/3	100	100	100	5	90	185	online	online	blacklisted	blacklisted	blacklisted
INFN-PAVIA	T3	IT	ACTIVE	100	2/2	no-test	no-test	100	no.data	100	no.data	online	online	online	online	online
INFN-ROMA1	T2D	IT	ACTIVE	100	3/3	100	100	92	99	86	228	online	online	online	online	online
INFN-ROMA2	T3	IT	ACTIVE	0	2/2	100	0	0	213	88	79	test	online	online	online	online
INFN-ROMA3	T3	IT	ACTIVE	100	2/2	no-test	100	91	43	80	29	online	online	online	online	online
INFN-T1	T1	IT	ACTIVE	99	2/2	100	100	91	646	99	52	online	online	blacklisted	online	blacklisted

SAAB results into action

- Starting from July 2013 SAAB is steadily in production for the automation of upload and write permissions
- More than 14k manual actions have been taken over
 - N.B.: whitelisting actions are performed only when appropriate before a blacklisting status expiration



Next: SAAB extension

- SAAB is going to be plugged in also with **Get** and **Del** SAM tests
- **Site-granularity** instead of DDM EP-granularity is being pondered over
 - in most of the cases when a site's DDM EP is blacklisted, other DDM EPs follow after a while
 - the idea is to blacklist other DDM EPs after the first if tests being failing, i.e. without waiting for the 0% threshold to be reached
- SAAB **principles** can be **extended** to metrics for **other resources**, e.g.:
 - Performance metrics for **networks** (e.g.: *PerfSonar*)
 - Experiment-specific functional tests for data **transfer**
- SAAB is not experiment-specific, it may be extended to other needs

Conclusions

SAAB design prescripts and implementation

- proved **successful**, within ADC Automation, to **enhance** the task of storage resources blacklisting upon **performance** by substituting
 - manual actions with **automatic** ones
 - human educated guesses with **non-ambiguous** and **reproducible** decision criteria
- while drawn to meet the needs of performance-based automatic handling of ADC storage areas, it is generally applicable to any resources tested by a metric
 - envision the possibility of deploying automatic, **performance-based** actions on ATLAS ADC **resources** by means of a single, **uniform** and clear **approach**

Thanks.

Questions?