



www.ucit.fr

Analyze-IT and Predict-IT: Statistical-learning Tools for HPC Optimization

- The job scheduler is the central point of your HPC infrastructure. Detailed analyses of its history logs can unveil valuable information about the system behavior.
- Our tools **Analyze-IT** and **Predict-IT** are able to process that information and provide you with metrics and predictions that **help to increase cluster production and profitability.**

Analyze-IT and Predict-IT work with the following job schedulers:



Torque PBS

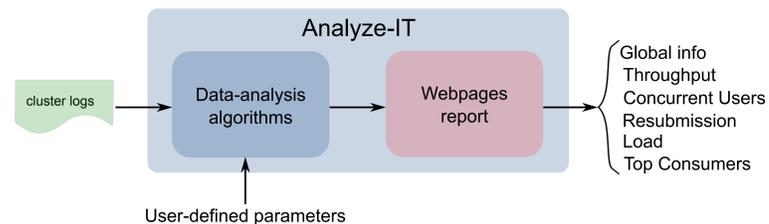
Analyze-IT

Understand
Advise
Optimize

- Goal: Understand cluster behavior in order to optimize it
- AIT extracts data from the cluster logs to reveal trends and patterns
- Key features
 - Overview of the cluster usage (cluster load, scheduler efficiency, job efficiency...)
 - Detailed metrics (job submission quality, cluster/nodes load, top 10 consumers...)

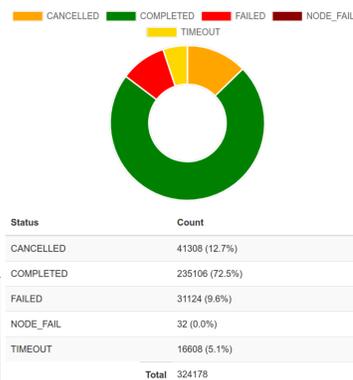
How it works:

- 1 Get data from job scheduler
- 2 Run Analyze-IT on the file (command line)



Analyze-IT allows you to:

- Filter data: by date, select only certain values in data fields (e.g., UID, JobNames, ...)
- Customize analysis: which analysis should be run, how many "consumers" should be displayed, time-based graphs resolutions, detect similar job names etc.
- Customize look and feel: colors, logo, analysis name...



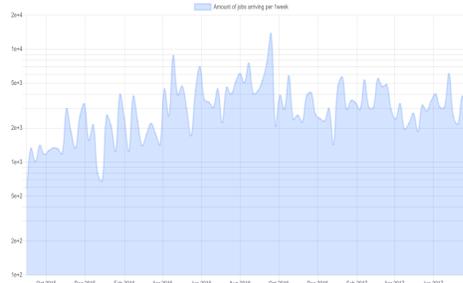
General metrics
Quickly visualize how many of the jobs turn out to be completed, cancelled, timeout, etc.

Identify Top Consumers
For example, which users have the highest number of allocated CPUs?

Optimization metrics
Data from the cluster logs are analyzed to generate quick assessment metrics.



Job arrival
Understand the historical pattern of job submission.



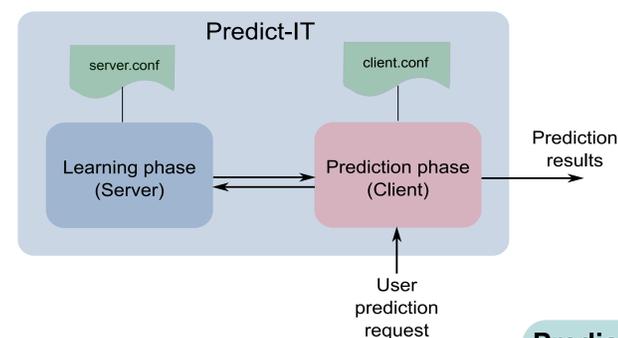
Predict-IT

Predict
Plan
Prescript

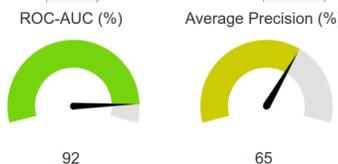
- Goal: Predict cluster behavior in order to enforce that submitted jobs will end up correctly.
- It uses machine-learning algorithms to devise a computational model of the cluster.
- Key features:
 - STATE prediction: detects the risk of a job finishing in timeout.
 - EXECTIME prediction: predicts the walltime that should be set by the user at the moment of submission to prevent job timeout.

How it works:

- 1 Activate server side to model the cluster
- 2 Submit prediction request from client side



Quality of STATE prediction
How good is the prediction of TIMEOUT jobs?

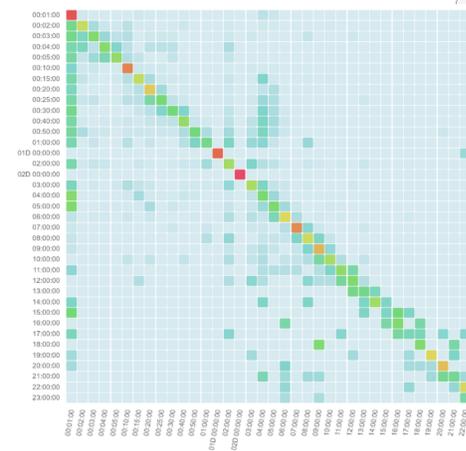


Predict-IT improves over time:
the more data you feed it, the higher the prediction accuracy can be.

Quality of EXECTIME prediction

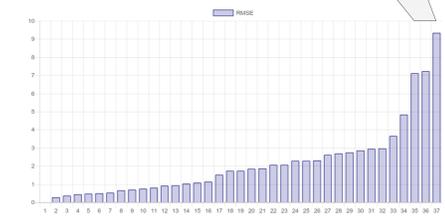
How good is the prediction of job execution time for each time bin (class)? The goal is to reach a perfect diagonal.

Orange and red squares represent a rate of correct classification of 80% or above.



EXECTIME prediction error

What is the level of error for each time bin (class)?



Feature importance
Which factors (UID, JobName, ...) carry more information about the STATE and EXECTIME of a job?

