

# Directory Search Performance Optimization of AMGA for the Belle II Experiment

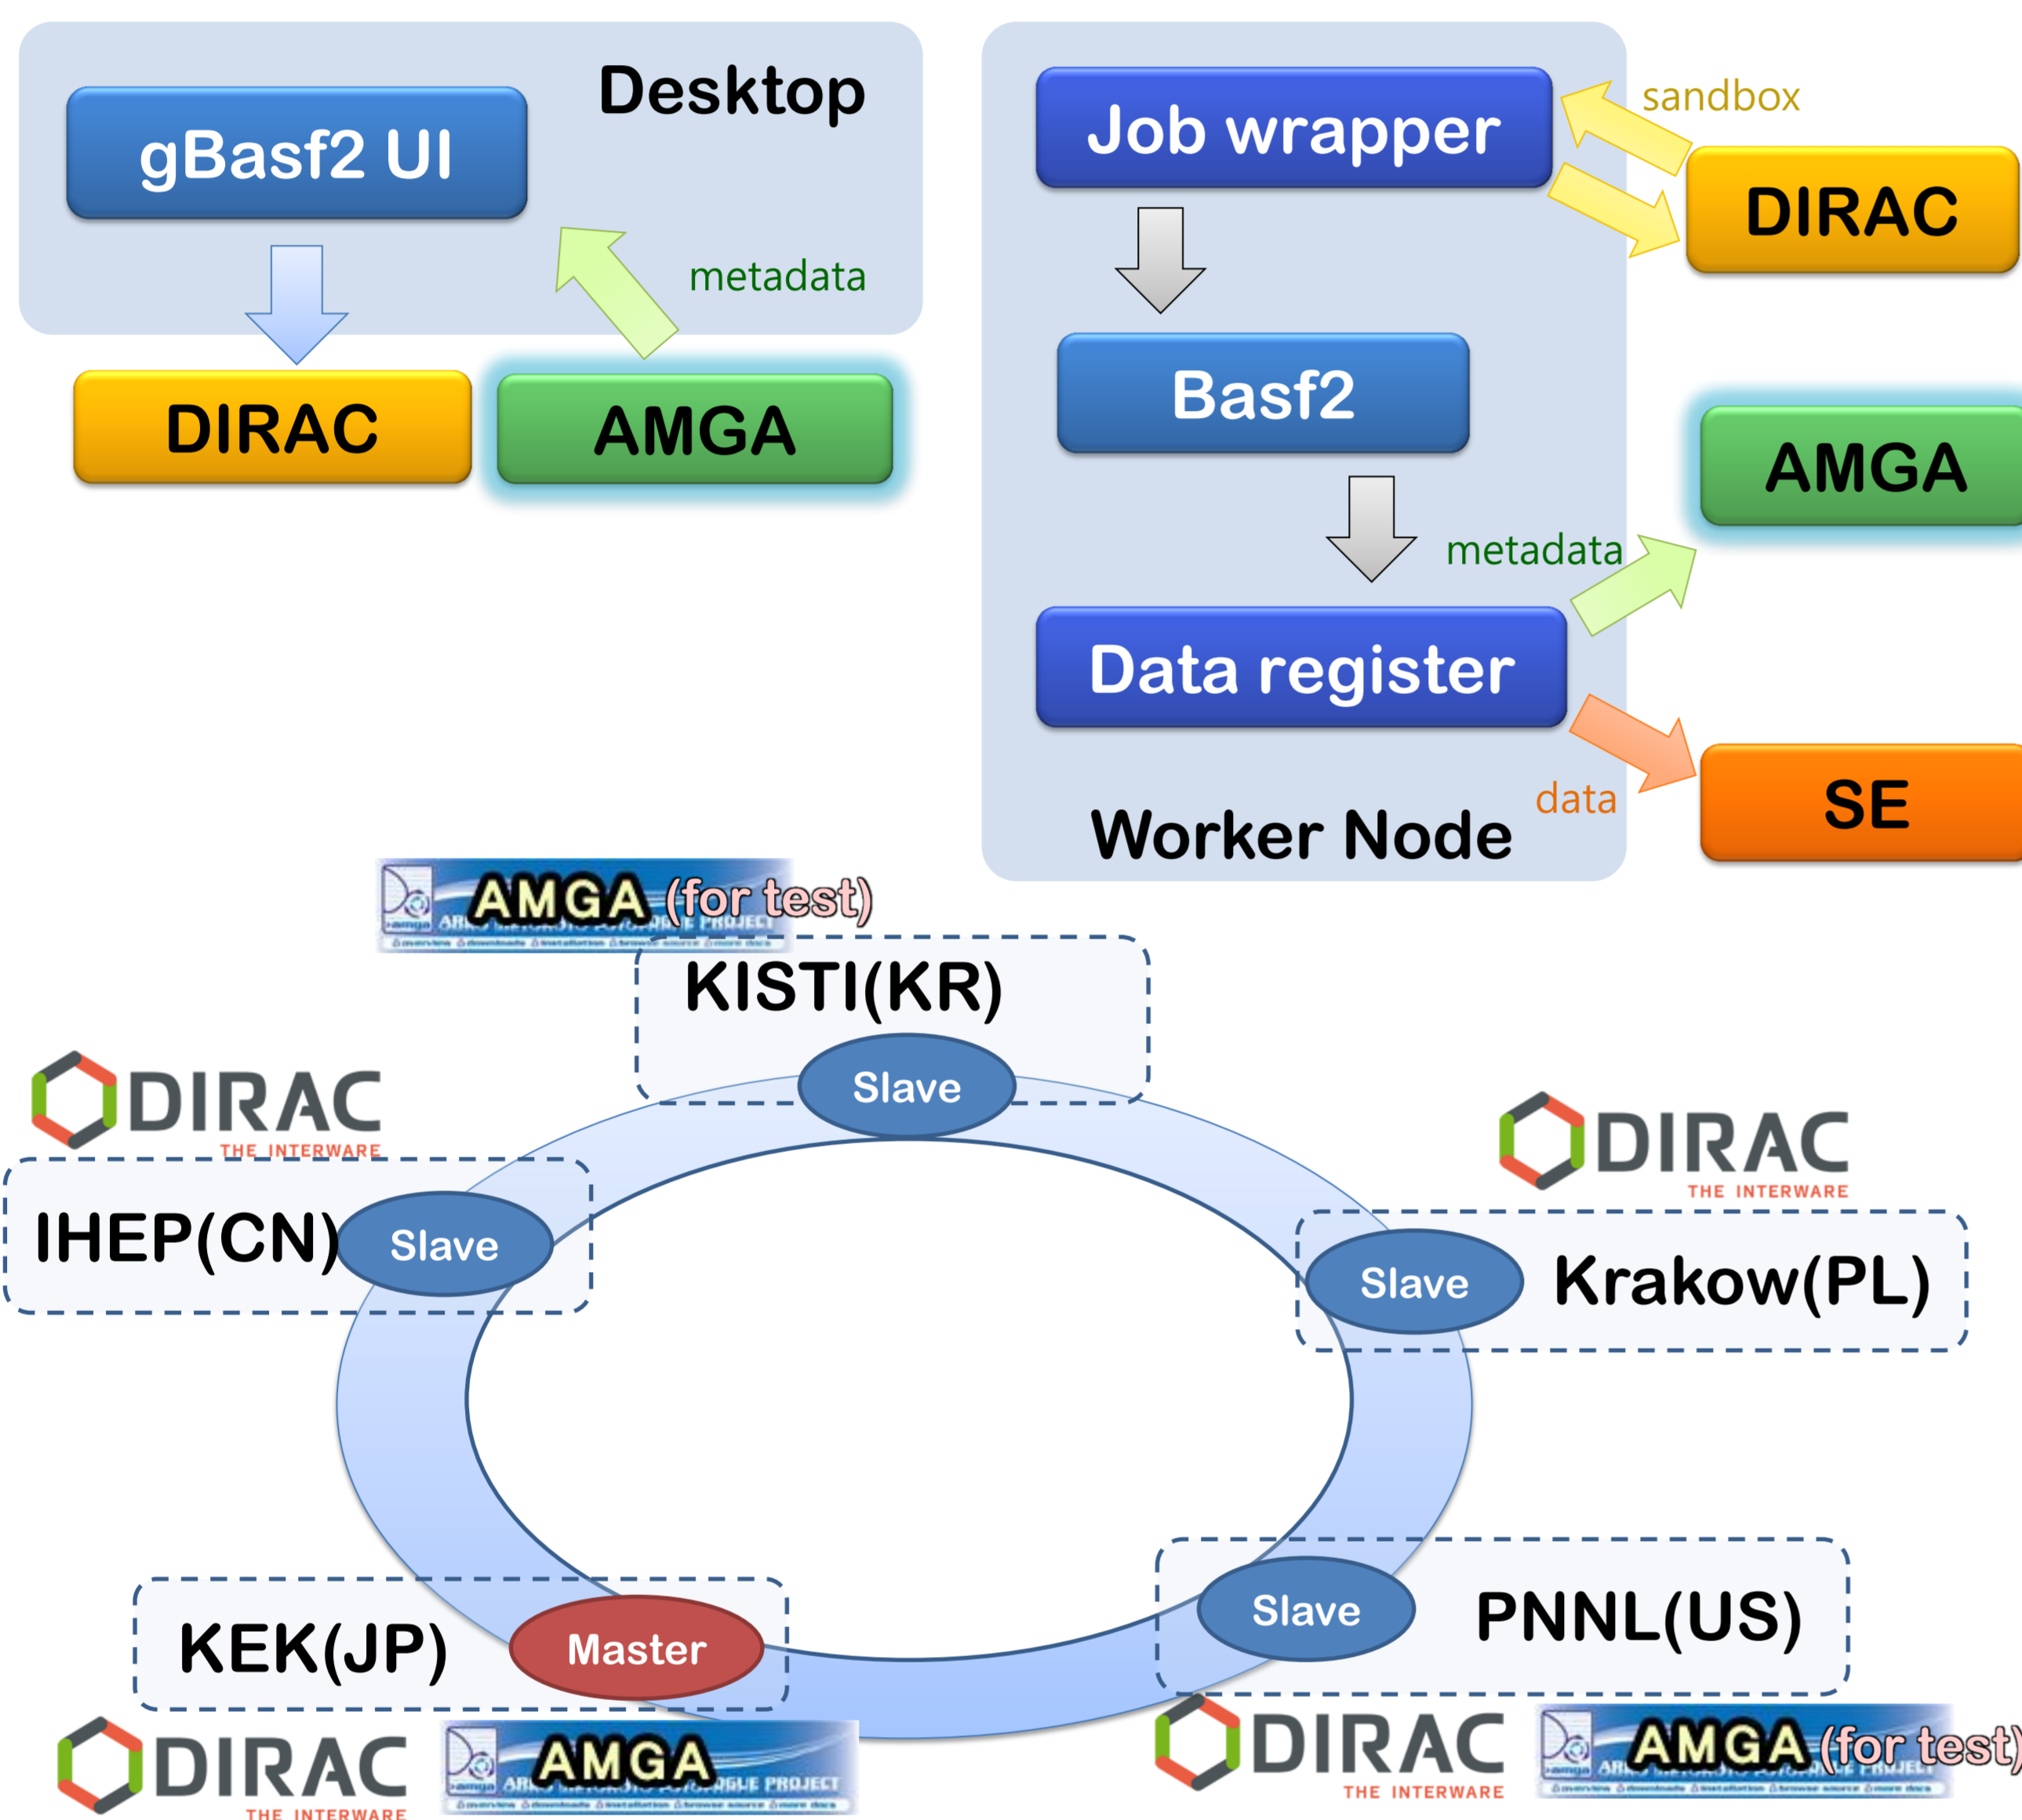
Geunchul Park, Jae-Hyuck Kwak, Taesang Huh, Soonwook Hwang

## AMGA (ARDA Metadata Grid Application)

URL: <http://amga.web.cern.ch/amga>

- An official EMI Metadata Software Product
  - included in the EMI Repository and EGI UMD (EGI Software Repository)
  - Registered at the XSEDE Technology Database(Evaluation competed!)
- provide and support a mechanism for locating scientific data files using descriptive information through Grid authentication
- Ensure competitive performance and scalability, along with a multi-threaded multi-process based on DB connection pooling, a hierarchical collection structure, replication and a federation mechanism
- KISTI: one of the product teams of EMI, contributing to the evolution and maintenance of AMGA
- EMI-3(Monte Bianco) AMGA(v2.4.0) Released(March, 2013)
- AMGA(v2.5.0) released to improve performance (for Belle II)

### AMGA in Belle II Distributed Computing System



## AMGA Performance Improvement

### Why too slow?

- All information of directories is stored in a masterindex table
- AMGA Uses a POSIX-style regular expression as a directory pattern matching method

```
Query > ls /belle/user/hideki
```

Query transition between AMGA and Back-end DB(PostgreSQL)

```
SELECT "directory", "table_name", "flags", "owner_name", "permissions", "acls" FROM public.masterindex WHERE "directory" ~ '^/belle/user/hideki/[^/]+$';
```

- But, the POSIX-style regular expression is costing too much

### Solution

- Avoid using the POSIX-style regular expression of the directory pattern matching method
  - Need to modify the masterindex table to use "low-cost" pattern matching style
- Add a new column to the masterindex table : parent
  - Keep a parent directory information
- Query can be simplified by using the parent column

```
SELECT "directory", "table_name", "flags", "owner_name", "permissions", "acls" FROM public.masterindex WHERE "parent" = '/belle/user/hideki';
```

## Test Result

### Test environment and command

- Virtual machine on the PC : Inter Core i7-4770K CPU
- VM Setting : 16G Memory, 2 Processors, 128GB HDD
- 'ls path' command used to test
  - Get 100 entries from target path
- Same schema such as a belle II metadata schema

```
* Schema
>> status : varchar(128) >> site : varchar(32) >> parentid : int
>> eventL : int >> runL : smallint >> date : timestamp
>> eventH : int >> events : int >> runH : smallint
>> stream : smallint >> versionid : smallint >> guid : varchar(32)
>> lfn : varchar(1024) >> experiment : int >> id : int
>> site : varchar(32) >> user : varchar(32) >> software : varchar(32)
```

### Test Result

| AMGA Version | 1 dir       | 10 dirs    | 50 dirs    | 100 dirs    | 500 dirs    | 1K dirs      |
|--------------|-------------|------------|------------|-------------|-------------|--------------|
| V2.4.0       | 100 entries | 1K entries | 5K entries | 10K entries | 50K entries | 100K entries |
| V2.5.0       | 0.0017      | 0.0018     | 0.0032     | 0.0038      | 0.0089      | 0.0139       |
| V2.5.0       | 0.0017      | 0.0017     | 0.0017     | 0.0017      | 0.0018      | 0.0018       |

| AMGA Version | 1K dirs      | 2K dirs      | 3K dirs      | 4K dirs      | 5K dirs      | 10K dirs   |
|--------------|--------------|--------------|--------------|--------------|--------------|------------|
| V2.4.0       | 100K entries | 200K entries | 300K entries | 400K entries | 500K entries | 1M entries |
| V2.5.0       | 0.0139       | 0.0247       | 0.0353       | 0.0454       | 0.0556       | 0.1081     |
| V2.5.0       | 0.0018       | 0.0018       | 0.0018       | 0.0018       | 0.0018       | 0.0018     |

| AMGA Version | 10K dirs   | 20K dirs   | 30K dirs   | 40K dirs   | 50K dirs   | 100K dirs   |
|--------------|------------|------------|------------|------------|------------|-------------|
| V2.4.0       | 1M entries | 2M entries | 3M entries | 4M entries | 5M entries | 10M entries |
| V2.5.0       | 0.1081     | 0.2069     | 0.3105     | 0.4116     | 0.5222     | 1.0268      |
| V2.5.0       | 0.0018     | 0.0019     | 0.0019     | 0.0019     | 0.0020     | 0.0026      |


AMGA Directory Search Performance Test  
On Virtual Machine

### AMGA V2.4.0

- The response time is increased at constant rate

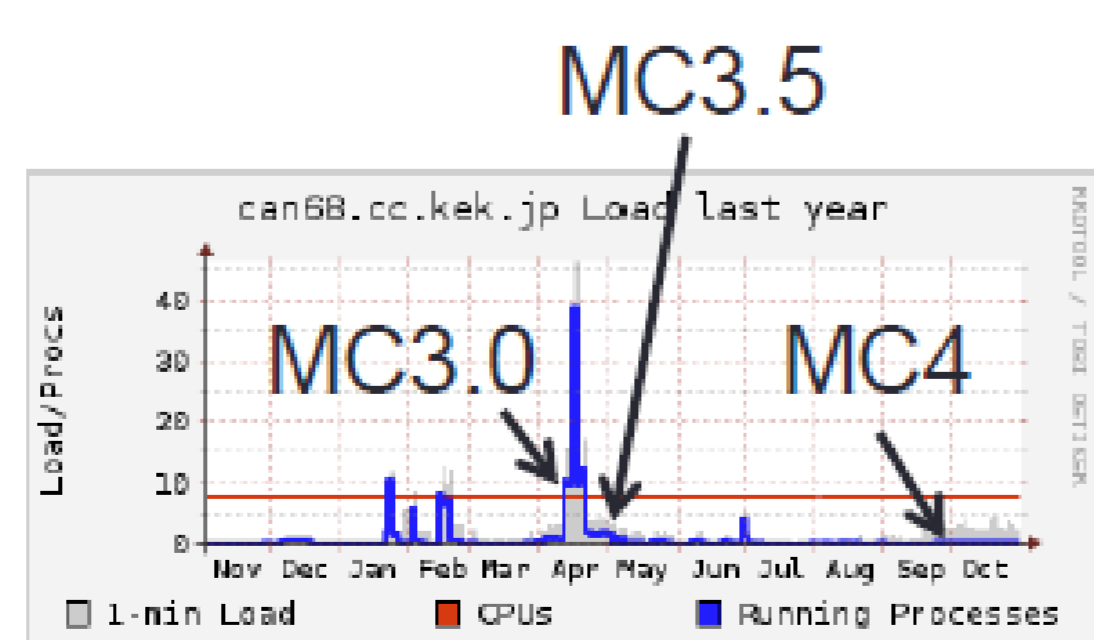
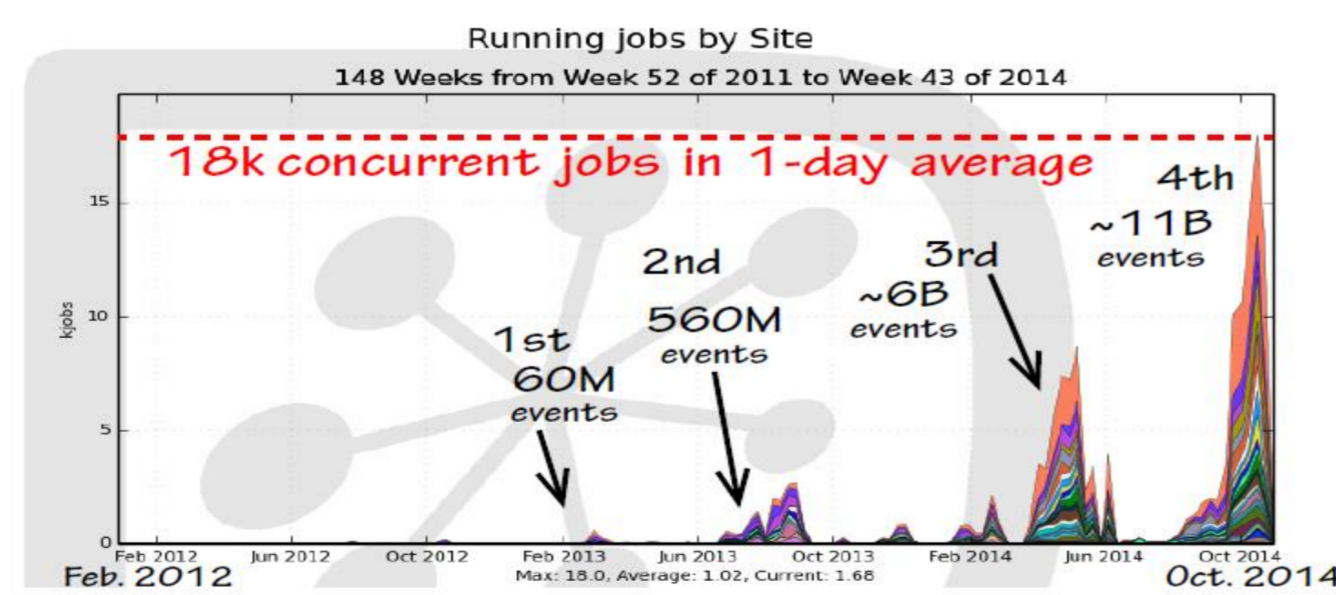
### AMGA V2.5.0

- The response time is stable at a constant value.



## Belle II MC Production Campaign

- Participated in the MC mass production campaign using Belle II software on Grid
  - Proves Belle II software architecture on the Belle II Grid
  - Observes what happens during massive resource consumption and shares the information
  - Grid sites, network, DIRAC, AMGA, and gbasf2 UI...
  - MC Production :
    - 15 countries/regions: Australia, Austria, Canada, Czech R., Germany, Italy, Japan, Korea, Poland, Russia, Slovenia, Taiwan, Turkey, Ukraine, USA
    - 31 sites: GRID, Cloud, local cluster is available
  - AMGA issue on MC 3
    - MC 3 : April 03 – May 02, 2014
    - Grown 2.5 times Jobs
    - Some AMGA issues reported
      - Access piled up and slow access
      - Heavy loads on the AMGA server
    - Checking AMGA performance
      - /w Hideki Miyake(KEK)
      - Checking performance: KEK AMGA Server VS KISTI AMGA Server
      - Identifying a cause : directory search performance
    - AMGA performance during MC4
      - AMGA 2.5.0 Version applied
      - CPU Load is very mild(vs MC3 and 3.5)



## Future Plans

- More tests for performance tuning and optimizing
  - Optimized postgresQL and AMGA parameters
  - Testing on the various postgresQL versions
  - Testing on the high-spec server
- Checking memory pile up issue
  - Enhancement of a sub-processor management algorithm
  - Add new option of refreshing sub-processors

