

# Use of R-GMA in BOSS

Henry Nebrensky (Brunel University)

VRVS 11 July 2003

Some slides stolen from various talks at EDG 2<sup>nd</sup> Review


(<http://documents.cern.ch/AGE/current/fullAgenda.php?ida=a021814>)

WP3 overview at GridPP middleware mtg. (???),

WP1-WP7, CERN, 18th June 2002


(<http://documents.cern.ch/AGE/current/fullAgenda.php?ida=a02943>)

# R-GMA




- Grid monitoring infrastructure
- Based on GGF GMA
- Discrete consumers and producers
- Registry acts as matchmaker

# R-GMA



**R-GMA**



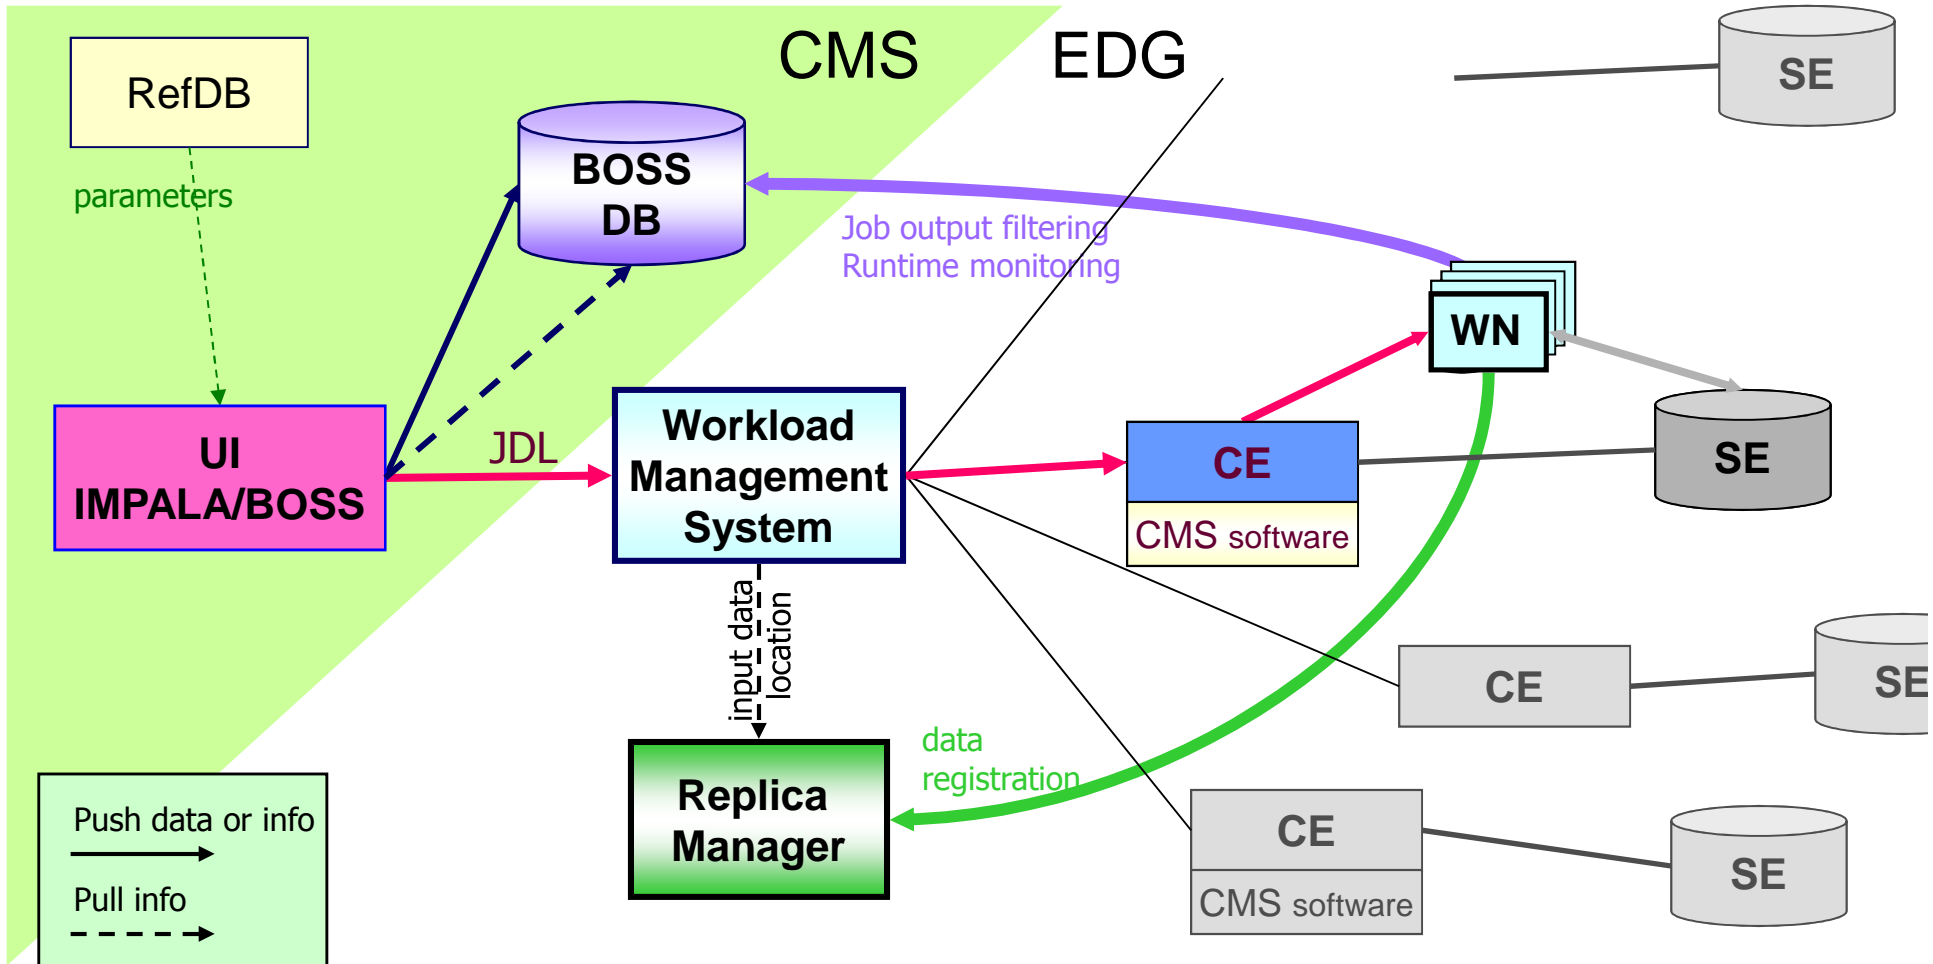
```
graph TD; Producer[Producer] -- subscribe --> Registry[Registry]; Consumer[Consumer] -- lookup --> Registry; Producer --> Consumer;
```

- Use the GMA from GGF
- A relational implementation
- Applied to both information and monitoring
- **Creates impression that you have one RDBMS per VO**

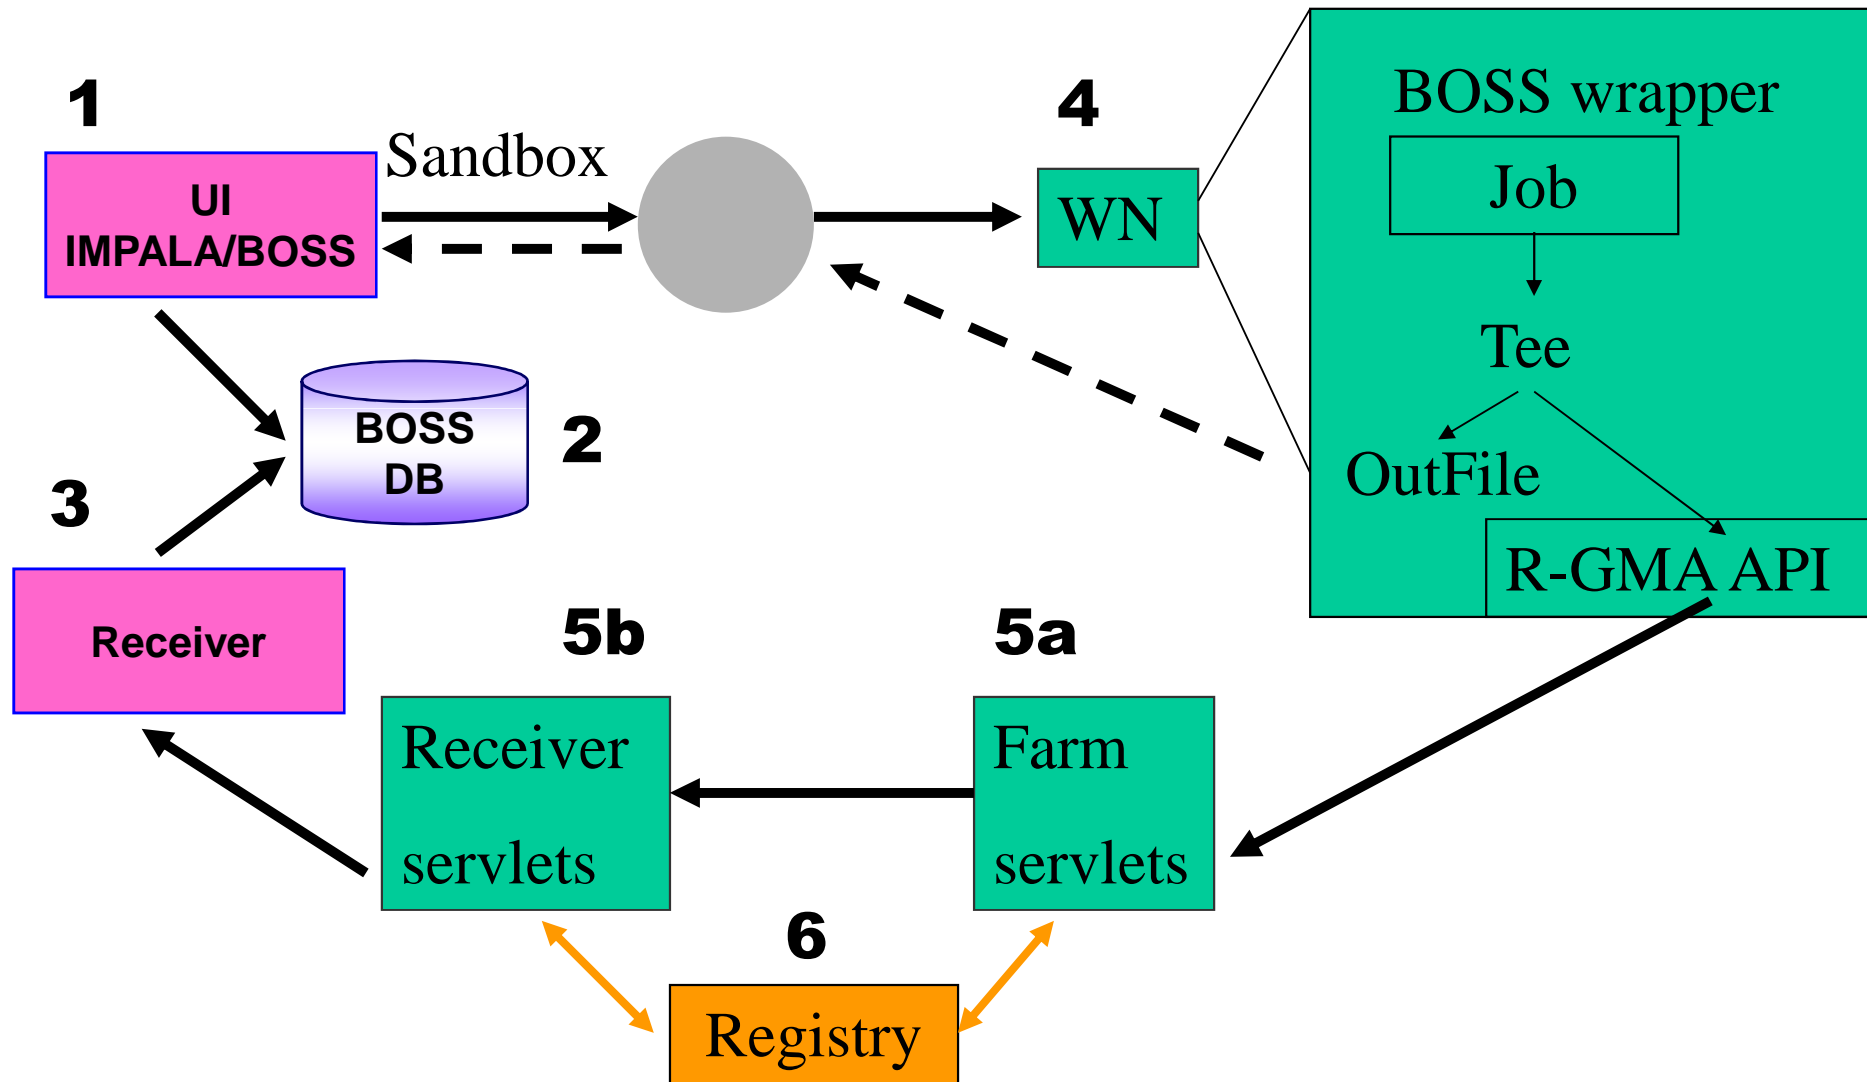
R-GMA and WP7 Steve Fisher/RAL - 12/6/2002 1

More on R-GMA see e.g. “**RGMA deployment**” at [http://www.gridpp.ac.uk/gridpp7/gridpp7\\_fisher.ppt](http://www.gridpp.ac.uk/gridpp7/gridpp7_fisher.ppt)

# BOSS



# Use of R-GMA in BOSS



# Use of R-GMA in BOSS

- Publish each update into R-GMA as a separate message – separate row
- Each producer gives host and name of “home” BOSS DB, and jobId; this identifies it uniquely
- Receiver looks for all rows relating to its DB; uses jobId and jobType to do SQL **UPDATE**

# Use of R-GMA in BOSS

The screenshot below shows the streamed output messages from a Brunel job (ID 112) being sent through R-GMA and displayed using the EDG Pulse tool from WP3. As Pulse can monitor multiple producers, it also shows the output from a longer job already running at Imperial (ID 72).

The receivers that update the BOSS databases use the `bossDatabaseHost` and `bossDatabaseName` fields to select only the relevant messages, so that the database at each institute is updated with only the information about its own jobs.

<http://www.brunel.ac.uk/~eestprh/GRIDPP/Index.htm>

SELECT * FROM bossJobExOutMessage						
bossDatabaseHost[]	bossDatabaseName[]	bossJobId[]	bossJobtype[]	bossVarName[]	bossVarValue[]	timeStamp[]
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	comment	I_am_fully_operational_and_all_my_circuits_are_functioning_perfectly.	1043425943
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	majorcount	204	1043425943
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	tick	15	1043425943
young.brunel.ac.uk:0	boss_v3_3_young	112	JOB	E_HOST	young	1043426585
young.brunel.ac.uk:0	boss_v3_3_young	112	JOB	E_PATH	/home/boss/boss-v3_3_pre5/CounterDemo	1043426585
young.brunel.ac.uk:0	boss_v3_3_young	112	JOB	E_USR	eesrjin	1043426585
young.brunel.ac.uk:0	boss_v3_3_young	112	JOB	T_START	1043426579	1043426585
young.brunel.ac.uk:0	boss_v3_3_young	112	counterdemo	comment	START...	1043426585
young.brunel.ac.uk:0	boss_v3_3_young	112	counterdemo	majorcount	0	1043426585
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	comment	Message_7:_This_is_message_number_7._Message_7_ends.	1043425948
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	majorcount	207	1043425948
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	tick	6	1043425949
young.brunel.ac.uk:0	boss_v3_3_young	112	counterdemo	majorcount	0	1043426590
young.brunel.ac.uk:0	boss_v3_3_young	112	counterdemo	tick	1	1043426590
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	comment	Brain_the_size_of_a_planet_and_he_has_me_count_to_twenty!_Bah.	1043425954
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	majorcount	209	1043425954
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	tick	17	1043425954
young.brunel.ac.uk:0	boss_v3_3_young	112	counterdemo	comment	I'm_sorry_Dave,_I'm_afraid_I_can't_do_that.	1043426595
young.brunel.ac.uk:0	boss_v3_3_young	112	counterdemo	majorcount	2	1043426595
young.brunel.ac.uk:0	boss_v3_3_young	112	counterdemo	tick	13	1043426595
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	comment	There's_a_pain_in_the_diodes_all_the_way_up_my_left_side.	1043425959
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	majorcount	212	1043425959
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	tick	8	1043425959
young.brunel.ac.uk:0	boss_v3_3_young	112	JOB	RET_CODE	0	1043426600
young.brunel.ac.uk:0	boss_v3_3_young	112	JOB	T_STAT	0.07s user 0.01s sys	1043426600
young.brunel.ac.uk:0	boss_v3_3_young	112	JOB	T_STOP	1043426600	1043426600
young.brunel.ac.uk:0	boss_v3_3_young	112	counterdemo	comment	That's_all_folks!	1043426600
young.brunel.ac.uk:0	boss_v3_3_young	112	counterdemo	majorcount	5	1043426600
young.brunel.ac.uk:0	boss_v3_3_young	112	counterdemo	tick	20	1043426600
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	majorcount	214	1043425964
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	tick	19	1043425964
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	majorcount	217	1043425969
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	tick	9	1043425969
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	comment	I'm_sorry_Dave,_I'm_afraid_I_can't_do_that.	1043425974
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	majorcount	220	1043425974
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	tick	20	1043425974
gw30.hep.ph.ic.ac.uk:0	boss_v3_3	72	counterdemo	majorcount	222	1043425979



# Use of R-GMA in BOSS (1)

- R-GMA smoothes “firewall” issues
- Consumer can watch many producers; producers can feed multiple consumers.
- Provides uniform access to range of monitoring data (WP7 network, etc.)
- Doesn't depend on other EDG components

## Use of R-GMA in BOSS (2)

- BOSS job wrapper uses an R-GMA StreamProducer and C++ API
  - Can define minimum retention period
  - No guarantees
- BOSS receiver implemented in Java
- Currently use R-GMA 3.1.39

# R-GMA

- Scalability not proven
  - GK a bottleneck?
  - Need separate CMS-specific R-GMA infrastructure?

# R-GMA

- WP3 large-scale testing:
  - “The Archiver the Boss job streams to can support 980 Stream Producers. The Stream Producer Servlet can support 420 Stream Producers. This means that you could not submit more than 420 Boss jobs to one site, if they had enough nodes to run all these on simultaneously. You could also not submit more than 980 Boss jobs across more than one site if you were using a single Archiver. Our Performance Analysis will hopefully improve these numbers.” *L. Field*

# Use of R-GMA in BOSS

- Augments BOSS journal files for real-time monitoring
- Need for large-scale testing now
- What are criteria for success ?
- How to deploy out there?

# Deployment

- **1UI 2DB 3RCVR** usually on same box? What instructions are needed?
- **4 WN**: have set of manual instructions that shouldn't clash with normal operation. LCFG?
- **5a 5b** are effectively EDG “MON” boxes. May only need one per site. Might need 8080 open inbound.
- **6 Registry**: do we use std. WP3 one or our own – this would flummox other R-GMA users e.g. WP7

<http://www.brunel.ac.uk/~eesrjjn/grid/grid.htm>