



INRDB the Internet Number Resource Database

Robert Kisteleki
Science Group Manager, RIPE NCC
robert@ripe.net



What is INRDB?

A system to store and retrieve long time series of Internet Number Resource related data, using reasonable computing resources.

Enables efficient access to across heterogeneous historical data.

Helps accomplishing RIPE NCC strategic goals:

- Trusted source of data
- Resource lifecycle management



Development goals

Design goal:

- Support the RIPE NCC's research and analysis efforts:
 - Access historical data about Internet Number Resources
 - Preparation for serving the data can be slow, retrieval should be quick and easy
 - Support various applications that use large amounts of data
- Store as much history as possible
- Provide a **single interface** for different datasets



Development results

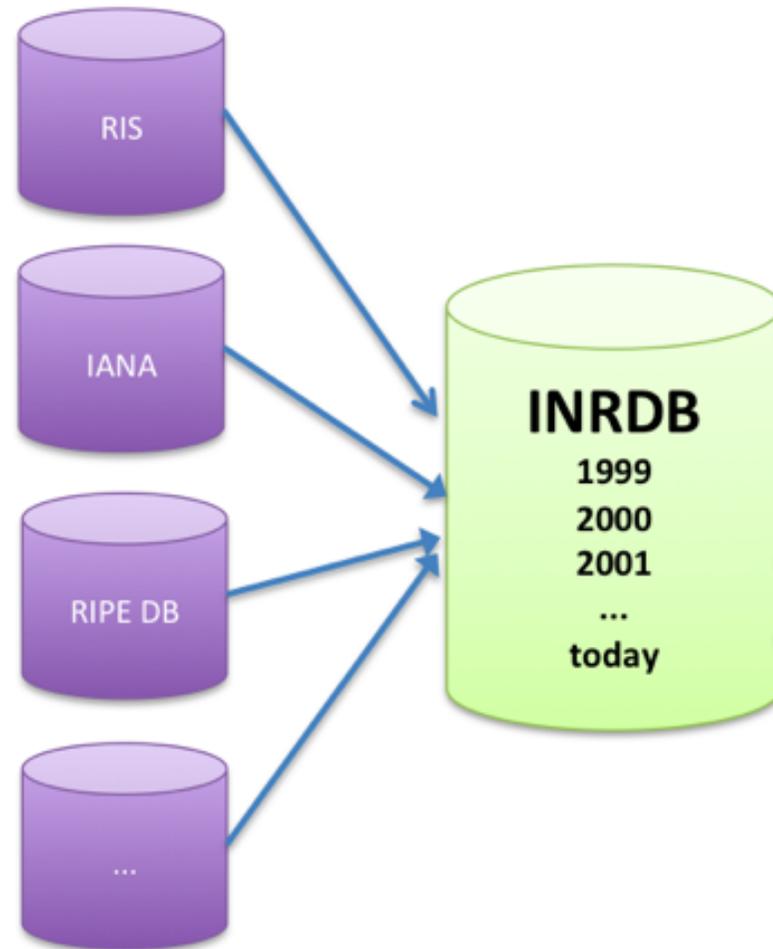
Results:

- Architecture is optimized for large databases
 - Think all of RIS table dumps and much more
- It works for us! :-)
- Recent evaluation: INRDB has high business value for the RIPE NCC, therefore steps are taken to turn INRDB into a production service.



INRDB overview

Input data sources



Applications





Concepts used by INRDB

Data stored/indexed:

- The “things” we observed = blobs
- Times when we saw those things = intervals
- Indexes exist for:
 - Resources (more/less specifics too, even if original DB does not have it)
 - Time intervals
 - Important non-numerical data



Concepts used by INRDB

Example:

```
BLOB: TABLE_DUMP2 | |B|200.219.130.11|28590|193.0.0.0/21|  
28590 12956 286 3333|IGP|200.219.130.32|0|0| |NAG| |
```

```
RES: 193.0.0.0/21
```

```
META: RIS_RIB, 200.219.130.11@rrc15
```

```
VALID: 2007-07-31T15:59:00Z - 2007-09-08T07:59:00Z
```



Currently available/served data sets

- RIPE NCC RIS table dumps (since 2000)
 - “normal” version: full RIB entries
 - “light” version: prefix + first transit AS + originating AS
 - “very light” version: prefix + originating AS
- All RIR statistics files (“delegations”)
- IANA assignment history
- Blacklists / spamlists: DROP and uceprotect
- GEOIP information from Maxmind



Currently available/served data sets

- Some CAIDA data sets
 - Reverse DNS lookups from Ark traces
 - AS relationships
- Various RIPE NCC internal databases

Some interesting numbers:

- ~160BN “input blobs” processed so far
- ~1.2BN blobs with ~8.5BN intervals stored / served currently
- We’re using 10 off the shelf servers



Background

INRDB is not a regular database:

- SQL was just too slow, too general for this
 - It's really difficult to store and index this much data
 - 100M+ records take just forever to index
- So we built a specialised storage and retrieval engine:
 - Geared towards storing blobs+intervals
 - Able to answer most frequent question types fast, while less common questions still in reasonable time



Background

INRDB peculiarities:

- It's not transaction oriented:
 - There's no such thing as “update” or “delete”
 - “Insert” is only effective for large number of items
- We have a separate processes:
 - “Update process” that crunches input data and produces INRDB “packages”
 - “Query process” that allows users to query these packages



Background

INRDB peculiarities:

- INRDB is effective for storing data that has moderate entropy
 - Routing tables tend to have lots of repetitions, constructing validity intervals makes sense here
 - Data sets that contain “random” measurements are really difficult to index effectively
 - But, we can store and index “event based” data too
- There’s nothing magical behind it
 - ~22K lines of C code



Architecture

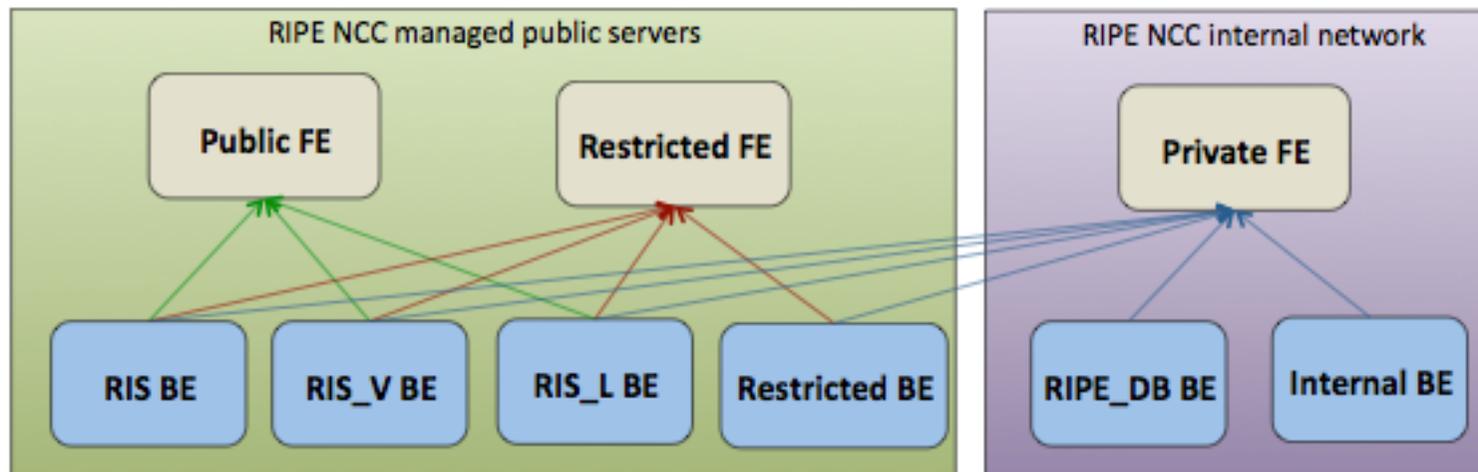
Architecture:

- Back Ends and Front Ends
- Use of multiple BEs and FEs enables load balancing
- Near linear scalability, packages/BEs can be added/removed any time
- Potential to include any time series about Internet number resources



Architecture

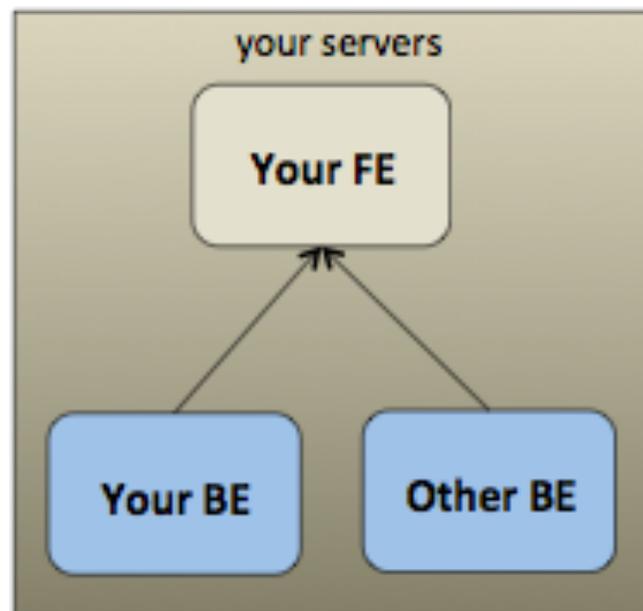
Current setup:





Architecture

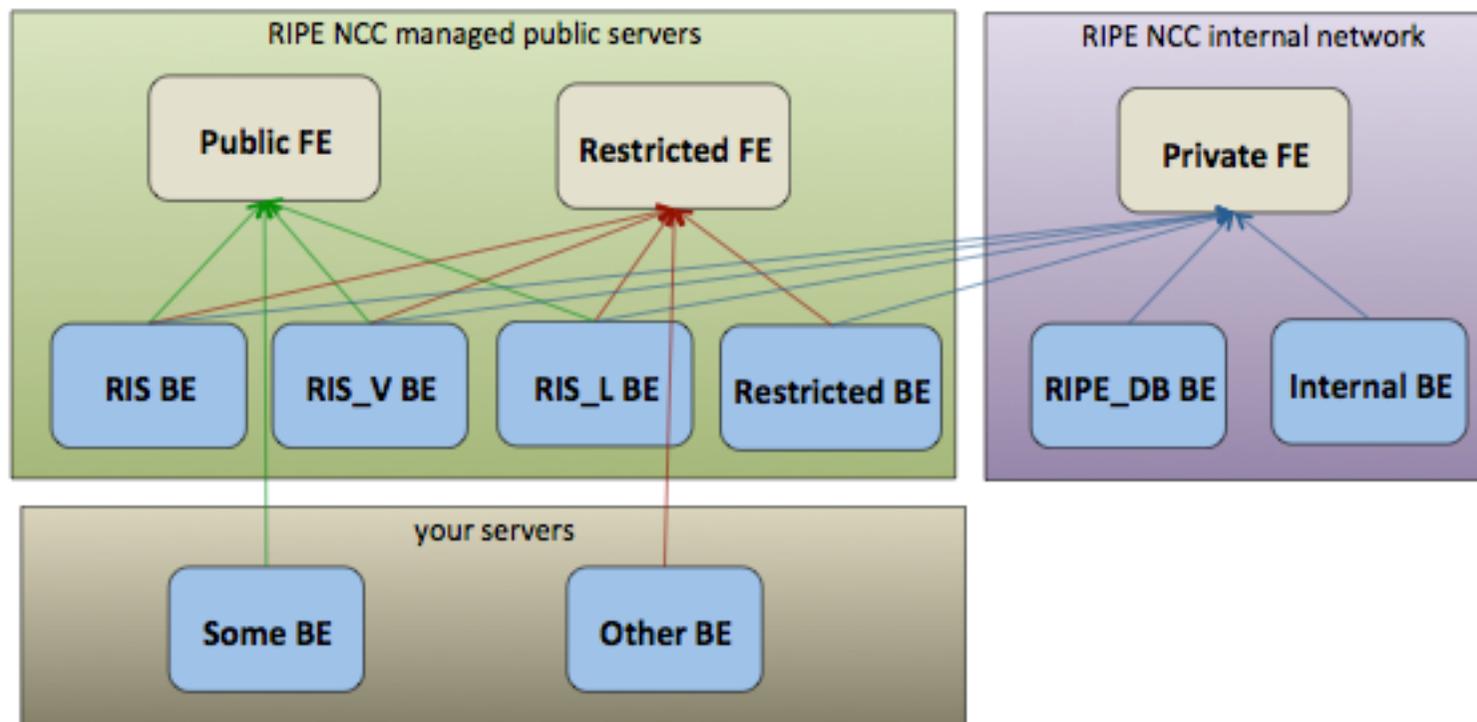
Building an independent setup:





Architecture

Extending the current setup:





Interfaces

We have developed a number of interfaces:

- “Raw” CLI access for quick checks and power users
- Perl and Java APIs
 - Object oriented access, most communication details are hidden
- JavaScript / JSON, XML, other interfaces are possible but we haven't built them



Query potential

Options (too many to list here, only examples):

- Restrict to time stamp / interval
- More/less specific searches on addresses (a'la RIPE DB but for all data)
- Non-numerical indexing
- Interval powers for RIS data (light / very light)
- Enable/disable report on:
 - Blobs, intervals, meta information, resources, powers, ...



Use of INRDB so far

- Structured analysis:
 - Membership demographics
- Ad-hoc analysis examples:
 - Mediterranean Cable Cuts
 - YouTube hijacking
- Prototype applications:
 - Registration Data Quality measurements (RDQ)
 - Resource EXplainer (REX)
- Numerous ad-hoc queries, quick checks, etc.



Summary

We managed to build a database that serves our needs:

- Stores data from a number of **different, large** data sets
- Provides a **uniform interface** for all this data
- Provides **indexing** on a number of properties
- Makes our research and analysis efforts **possible, or** at least much **easier** than before



Summary

It works for us, it may work for you!

- Other data sets can be plugged into the running system, provided they are run through the update process first
 - It doesn't matter *who* actually serves the data, the architecture can hide that
- We can also share the code with you, so you can play with it on your own.
 - There are no strings attached.



Demo

```

mc - /Volumes/disk01 - bash - 82x27 - #1
3333|IGP|193.232.244.111|0|0||NAG||
META: RIS_RIB, 193.232.244.111@rrc13
VALID: 2008-10-04T23:59:00Z - 2008-10-04T23:59:00Z

BLOB: TABLE_DUMP2||B|193.232.244.147|28917|193.0.0.0/21|28917 1299 3356 3333|IGP|1
93.232.244.147|0|0|28917:4000 28917:4001|NAG||
META: RIS_RIB, 193.232.244.147@rrc13
VALID: 2008-10-27T07:59:00Z - 2008-10-27T07:59:00Z

BLOB: TABLE_DUMP2||B|193.232.244.111|8331|193.0.0.0/21|8331 9002 9002 9002 12859 3
333|IGP|193.232.244.111|0|0||NAG||
META: RIS_RIB, 193.232.244.111@rrc13
VALID: 2008-11-21T15:59:00Z - 2008-11-21T15:59:00Z
VALID: 2009-02-04T15:59:00Z - 2009-02-14T15:59:00Z
VALID: 2009-02-18T16:00:00Z - 2009-03-07T15:59:00Z

BLOB: TABLE_DUMP2||B|193.232.244.187|9002|193.0.0.0/21|9002 12859 3333|IGP|193.232
.244.187|0|0||NAG||
META: RIS_RIB, 193.232.244.187@rrc13
VALID: 2008-11-21T15:59:00Z - 2008-11-21T15:59:00Z
VALID: 2009-02-04T15:59:00Z - 2009-02-14T15:59:00Z
VALID: 2009-02-18T16:00:00Z - 2009-03-14T15:59:00Z
VALID: 2009-04-02T23:59:00Z - 2009-04-03T07:59:00Z
VALID: 2009-07-08T07:59:00Z - 2009-07-08T07:59:00Z

BLOB: TABLE_DUMP2||B|193.232.244.15|39821|193.0.0.0/21|39821 9002 12859 3333|IGP|1
93.232.244.15|0|0|9002:9002 9002:64679|NAG||
  
```



Questions?

