



Sparta's Dual Kingship and PostgreSQL active-active

Boriss Mejías


Holistic System Software Engineer

Air Guitar Player

23rd of October, 2024 – PgConf EU



image from wikipedia: Sparta



Sparta's Dual Kingship and PostgreSQL Active-Active Mirage




image from wikipedia: Sparta

How can we deploy active-active in PostgreSQL?



It depends



~~It depends~~
You don't need active-active



It depends



Why do you need active-active?
What is actually the problem?



Scale-out

High Availability

Geo Redundancy



“If you want to do history, you have to be committed to getting it right, rather than being right.”

*Michael Livingston
according to Myke Cole*



databases

“If you want to do history, you have to be committed to getting it right, rather than being right.”

*Michael Livingston
according to Myke Cole*



computer science
active-active
PostgreSQL
databases

“If you want to do history, you have to be committed to
getting it right, rather than being right.”

*Michael Livingston
according to Myke Cole*



origin

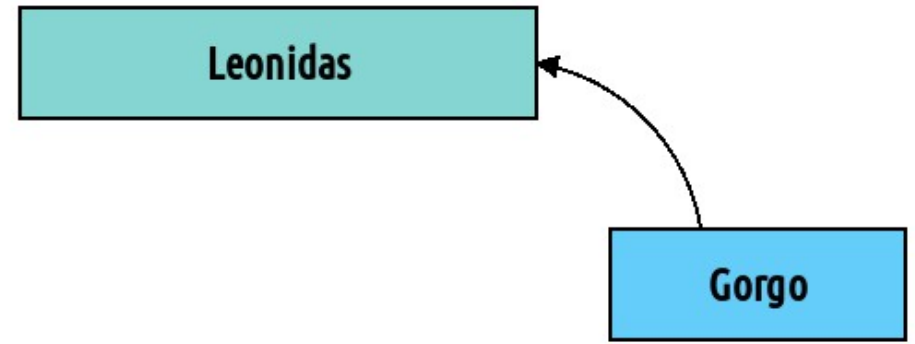
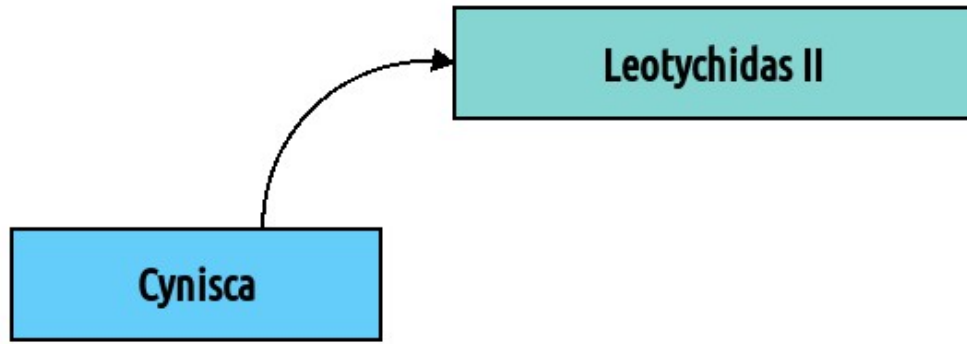




`origin=none`

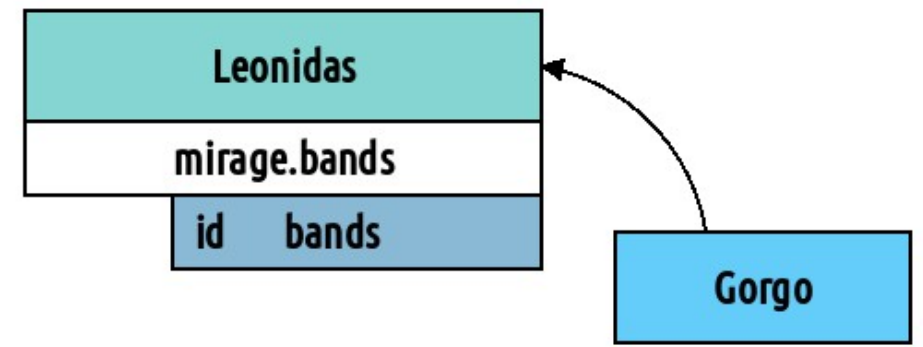




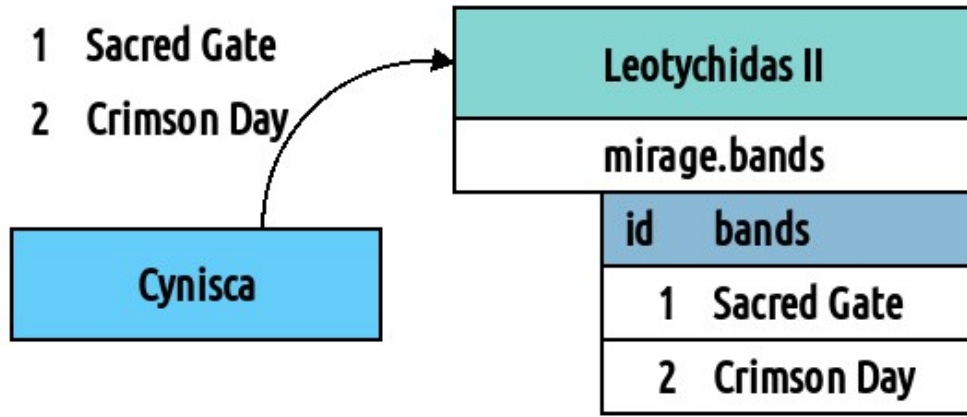
```
CREATE  
SUBSCRIPTION sub_eurypontid  
CONNECTION 'host=agis dbname=sparta user=sparta'  
PUBLICATION agiad  
WITH (origin=none);
```


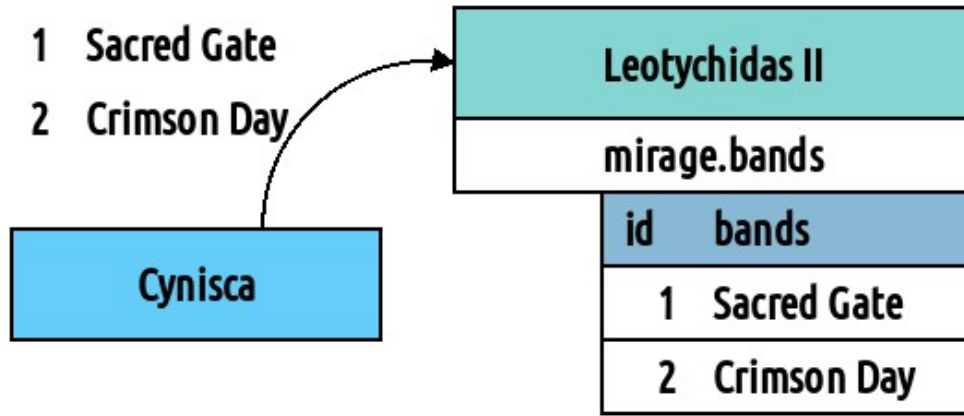


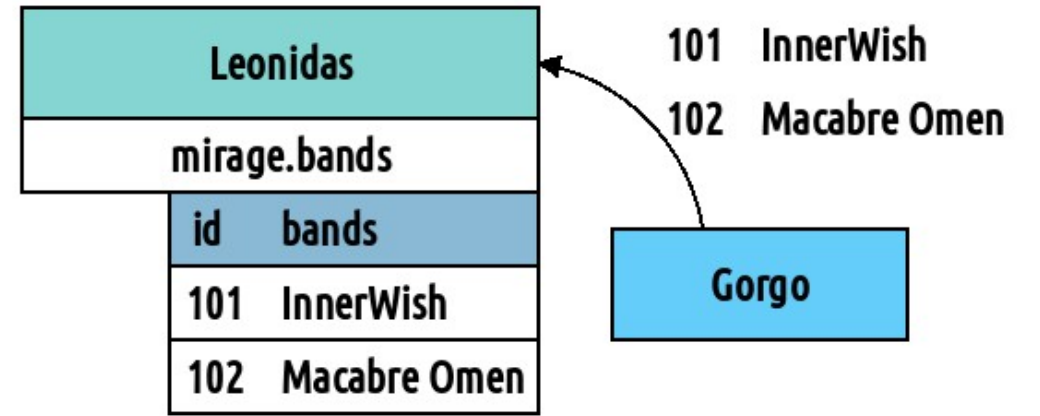






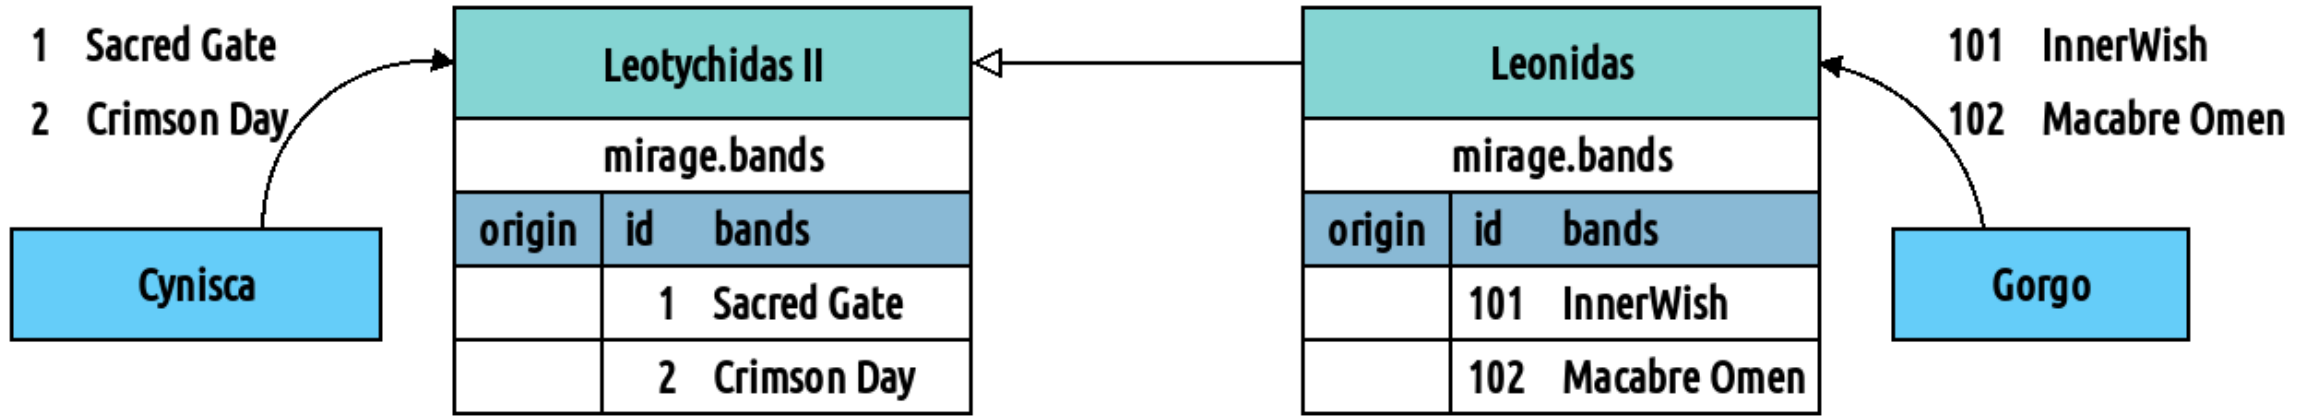


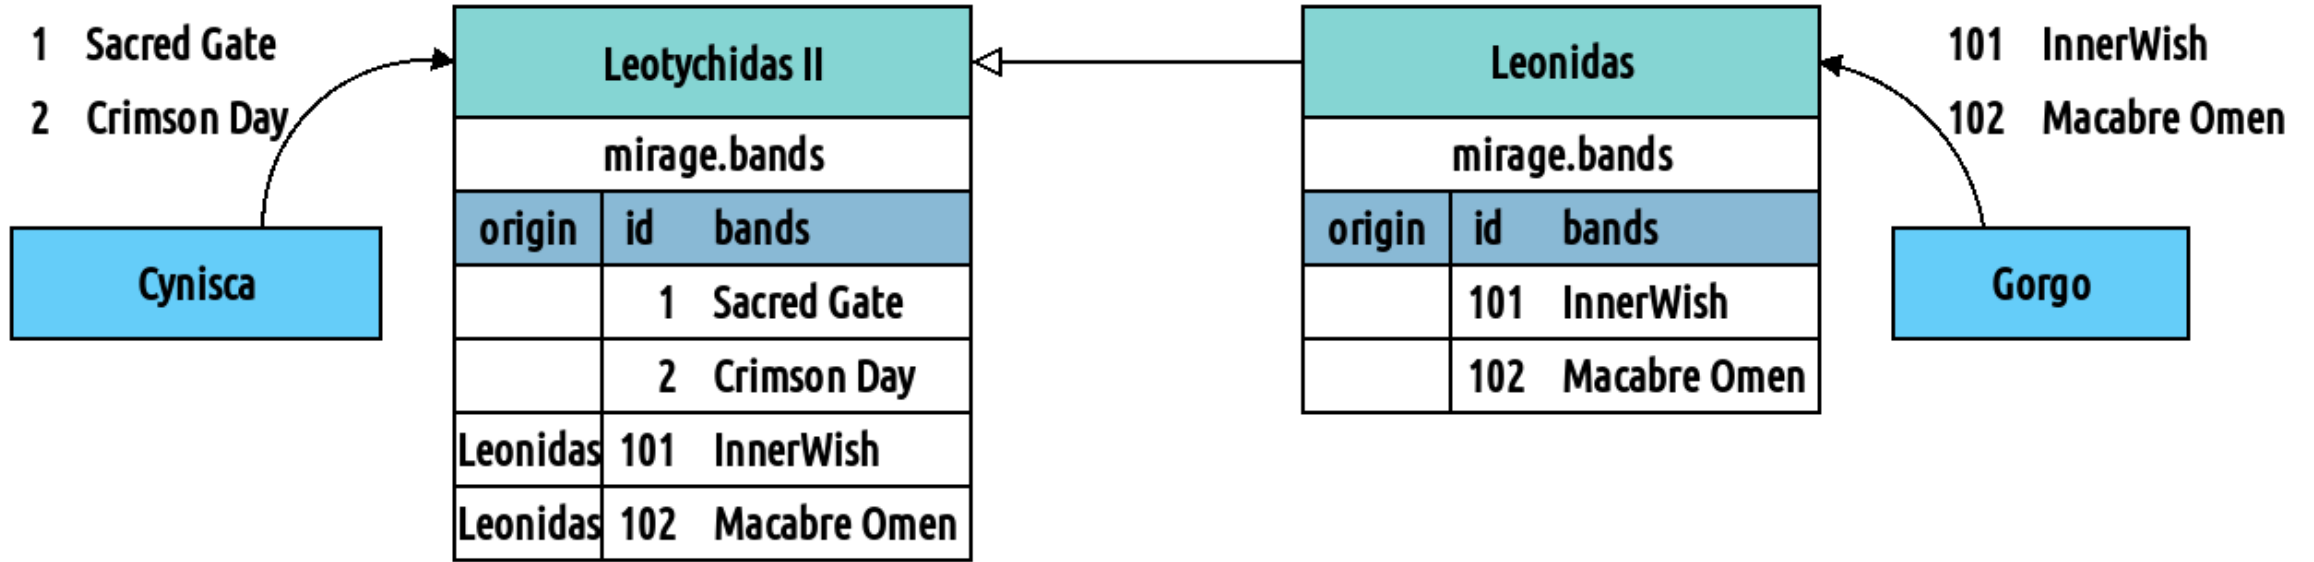


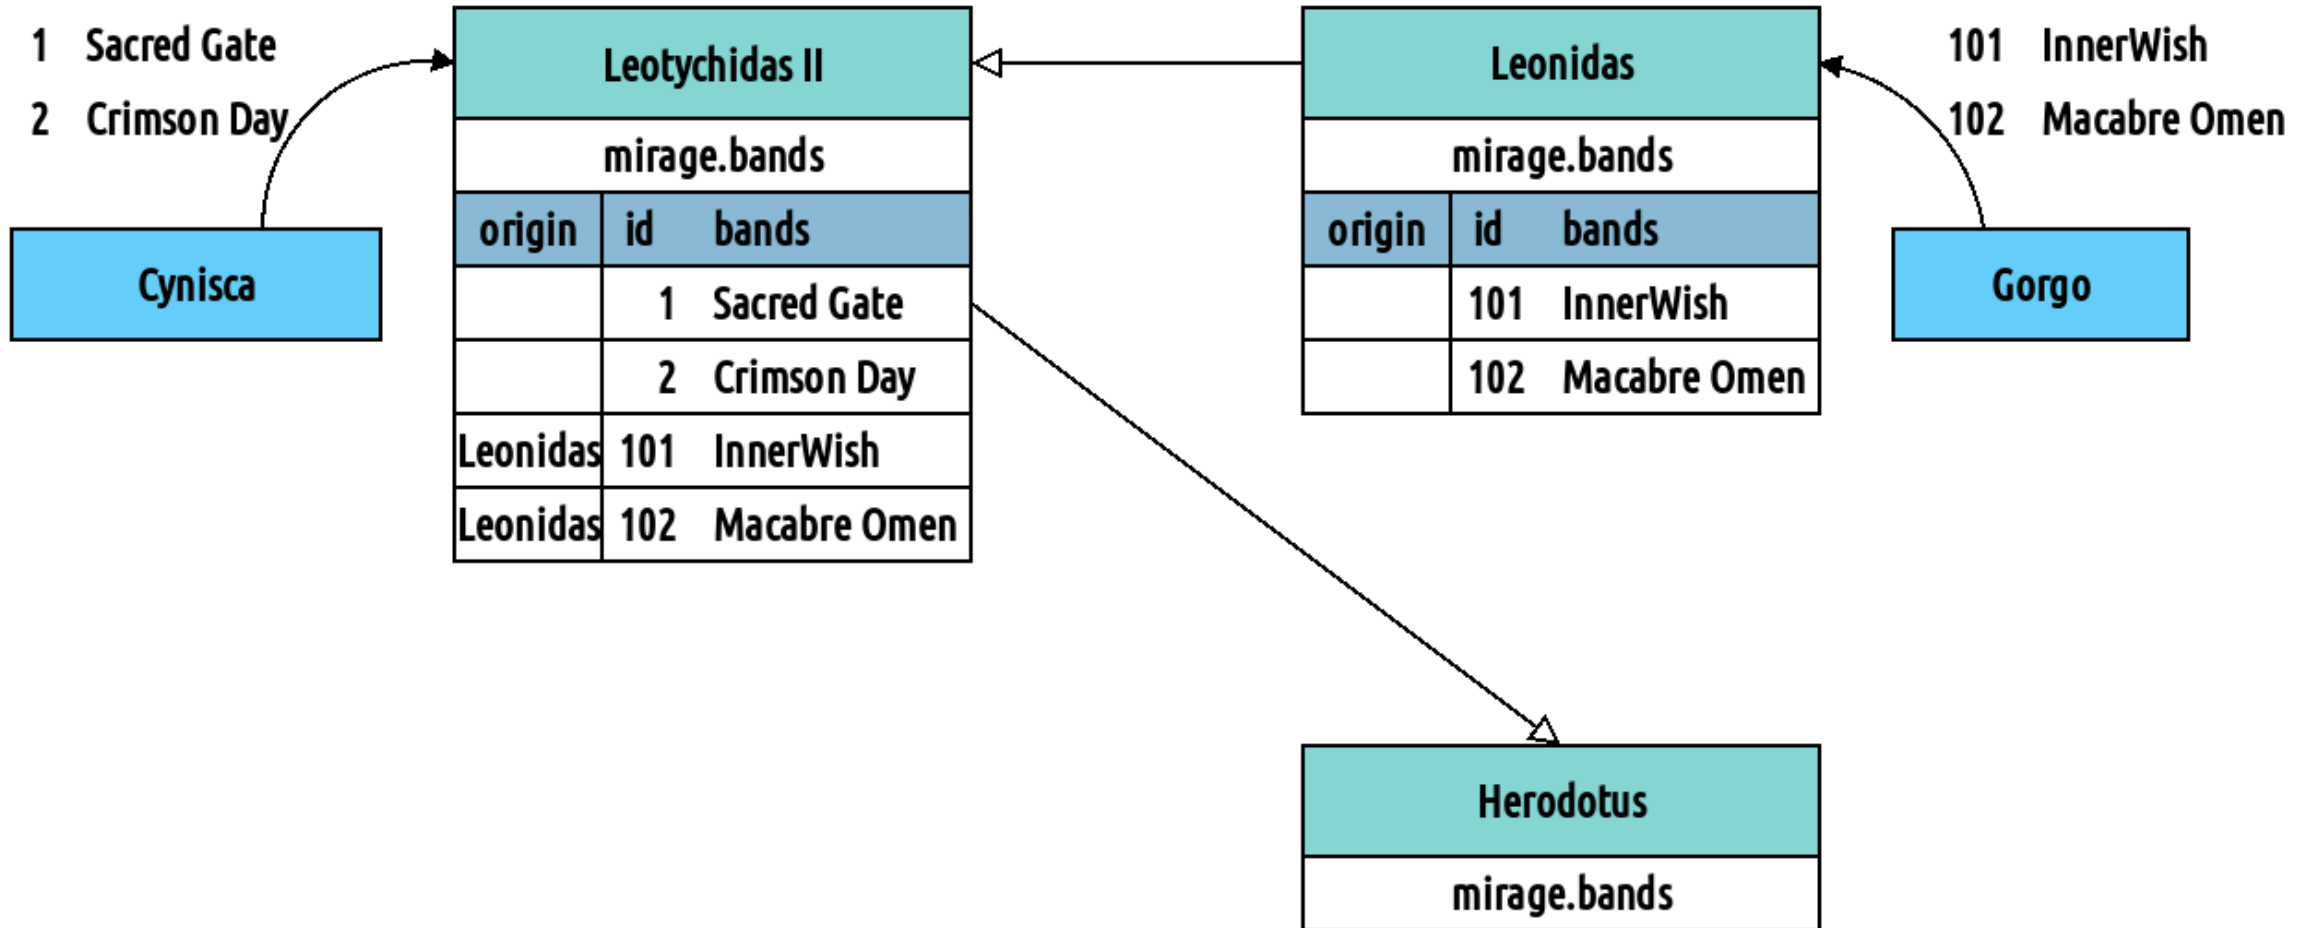


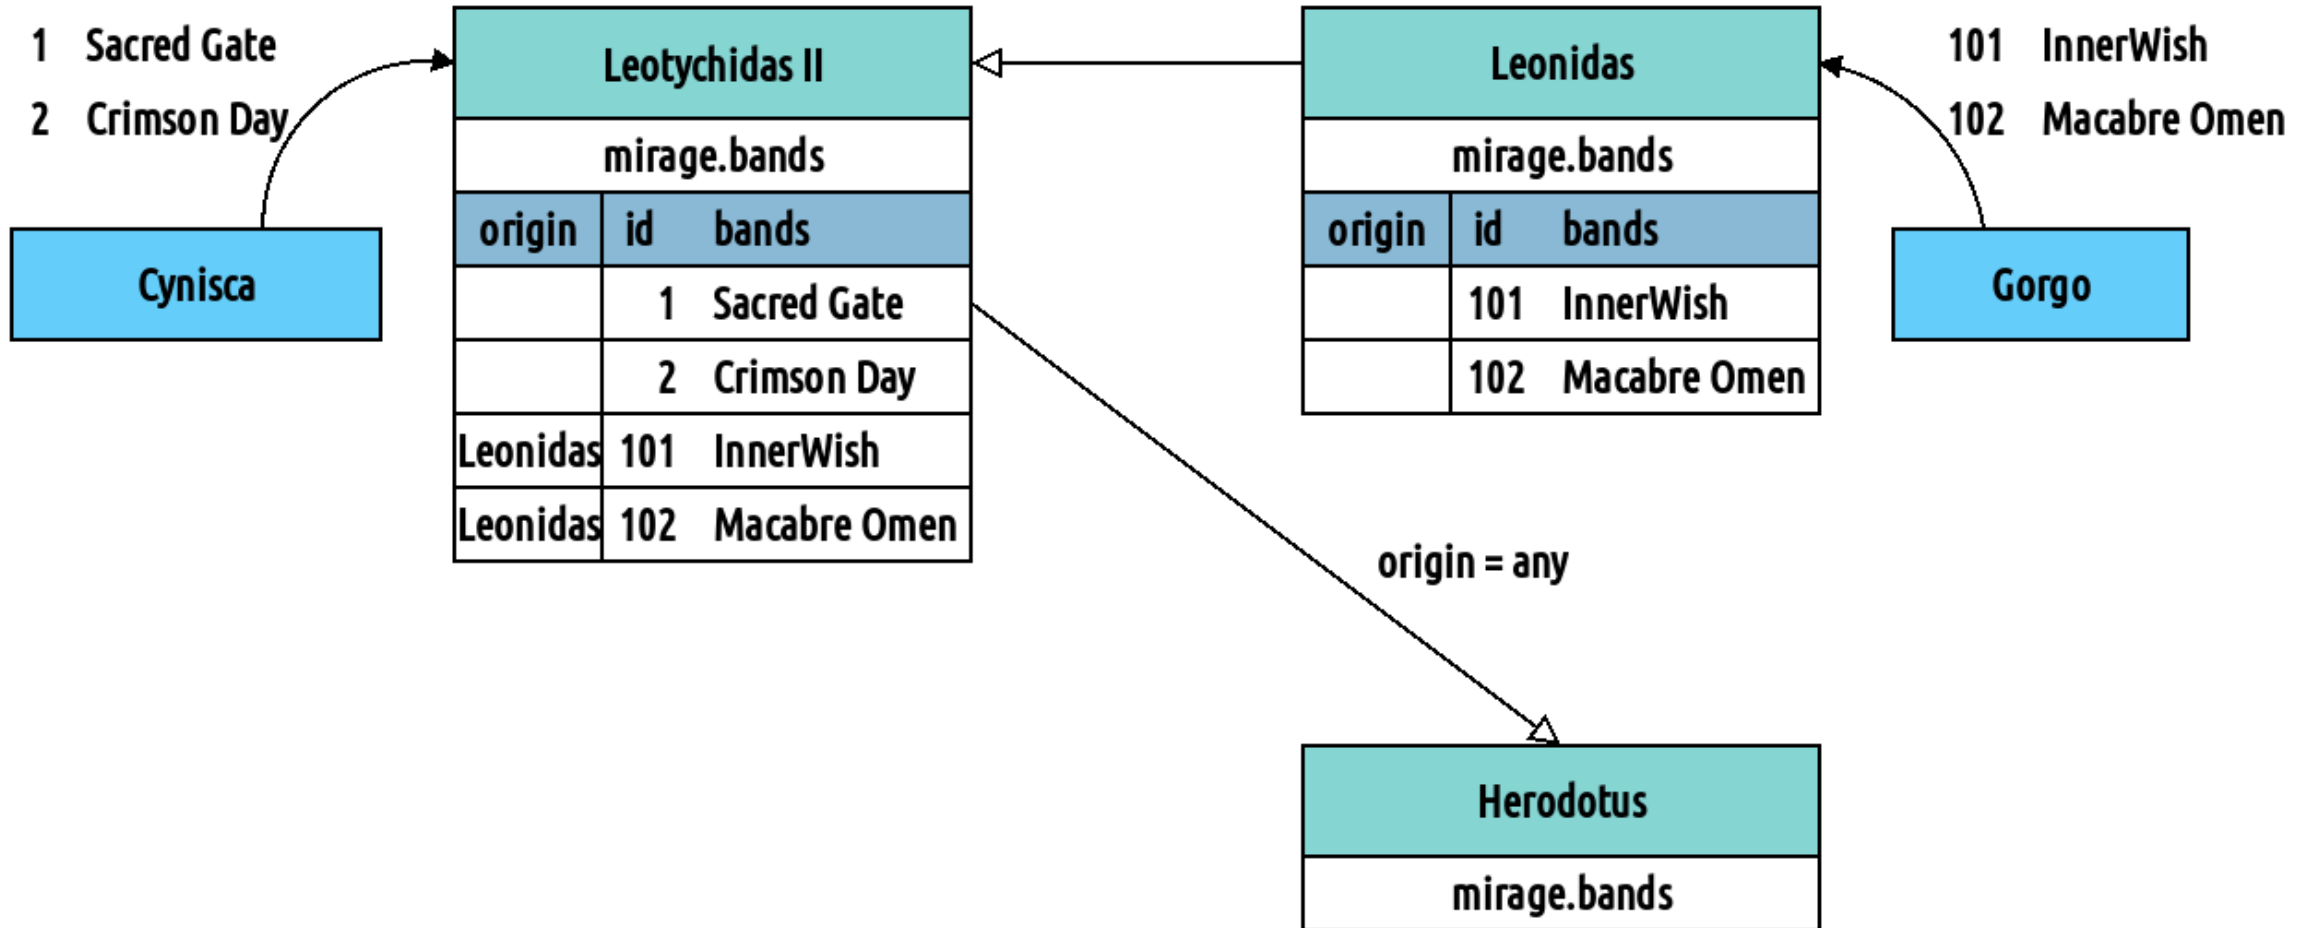


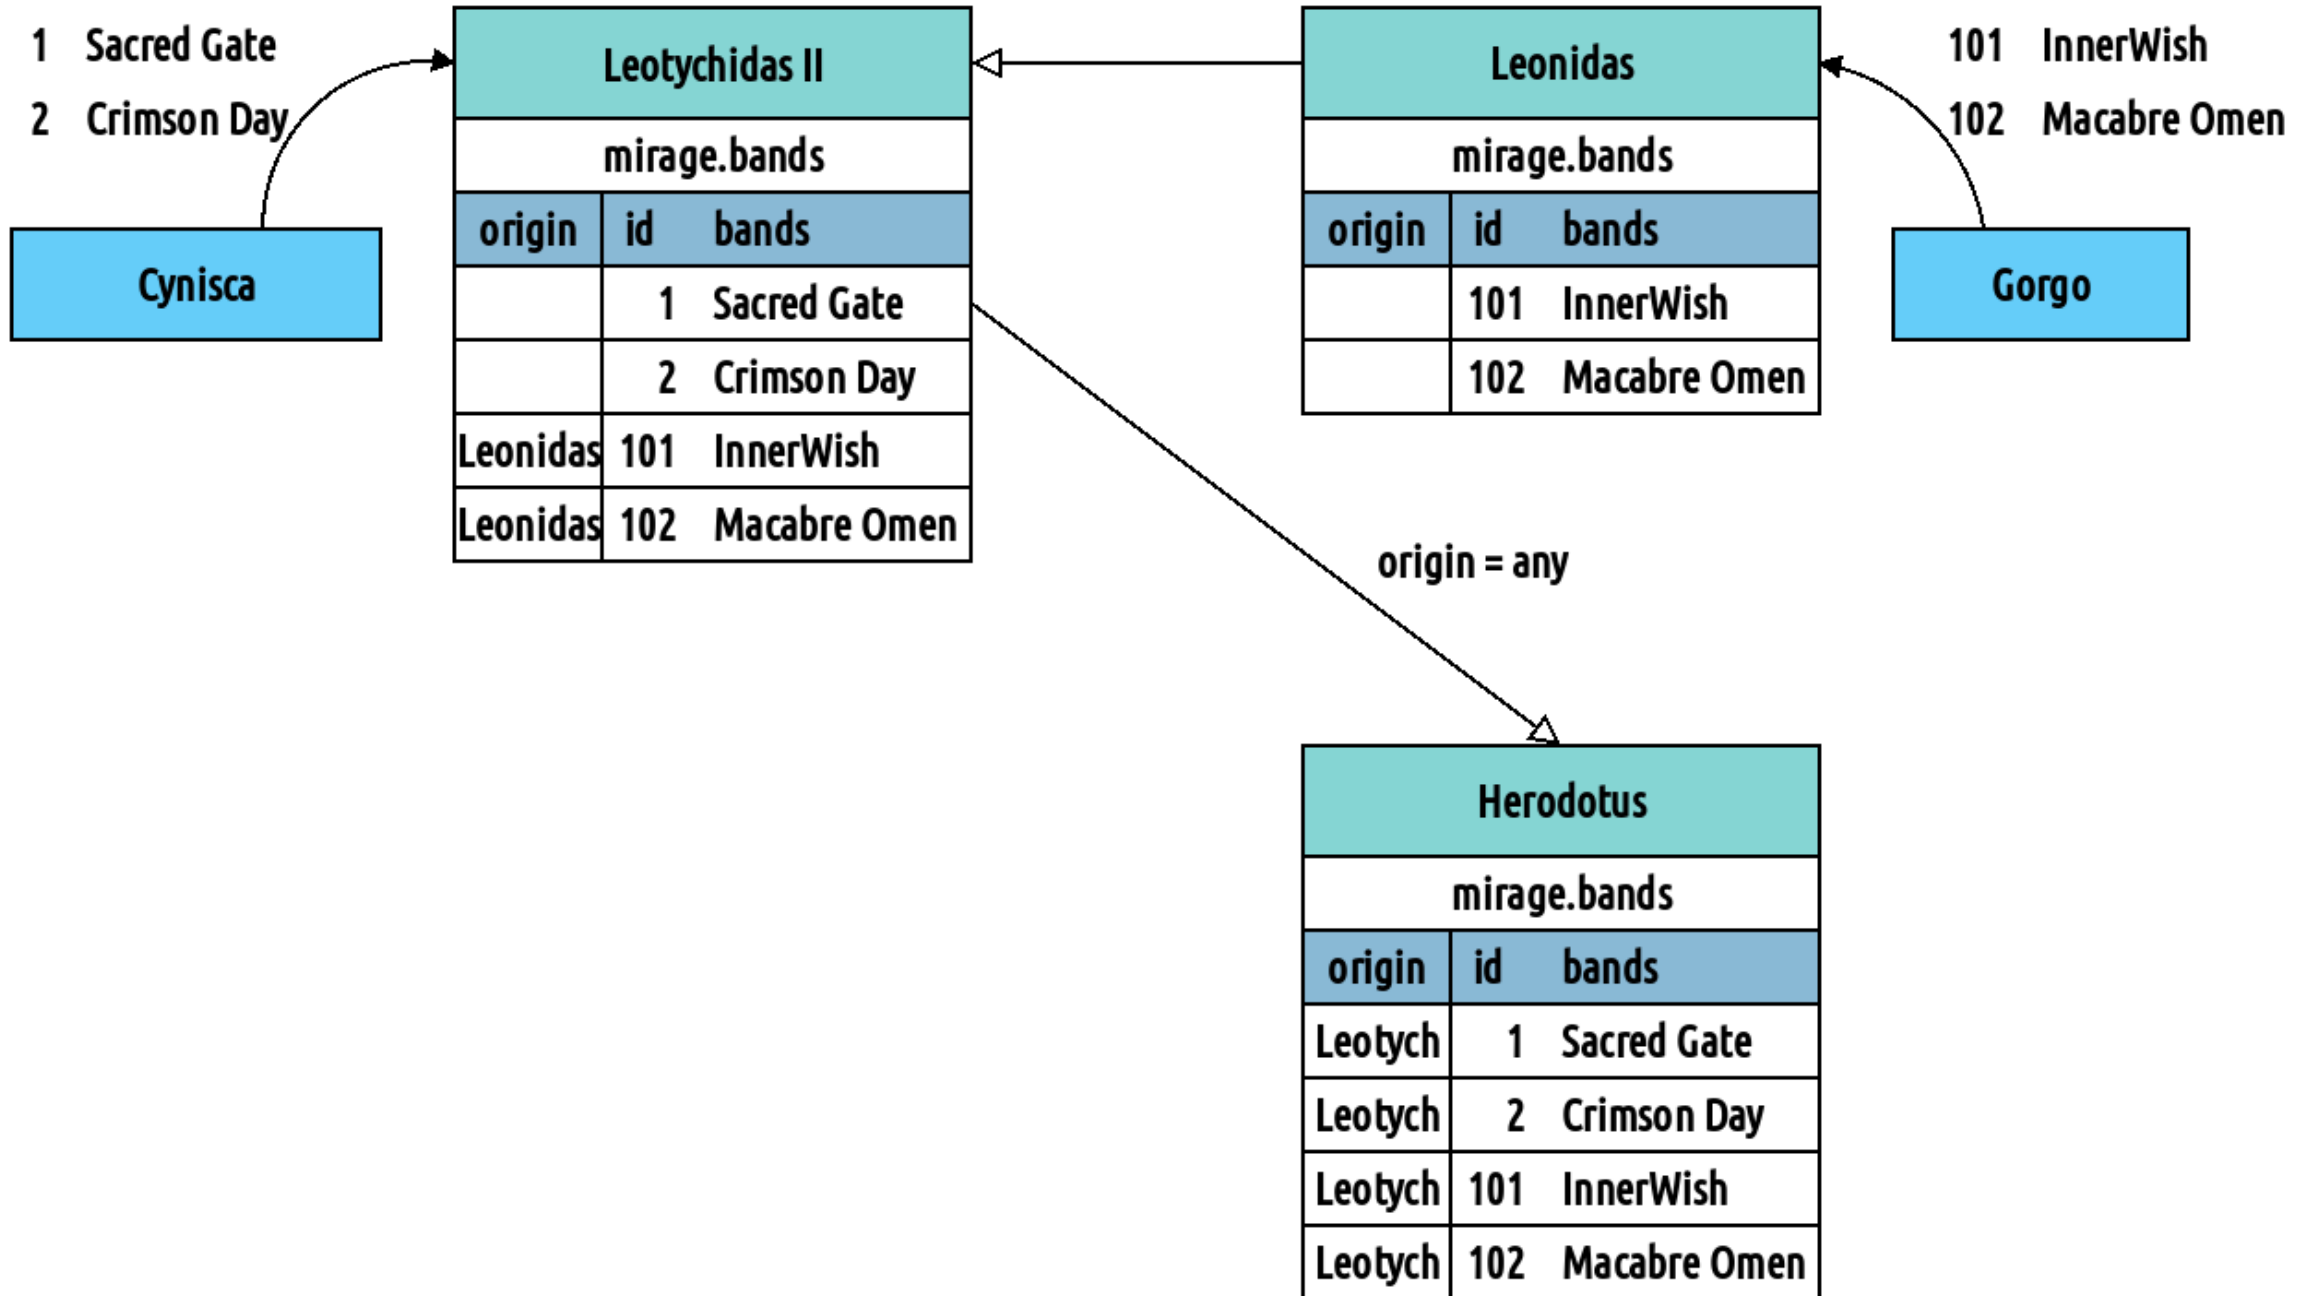


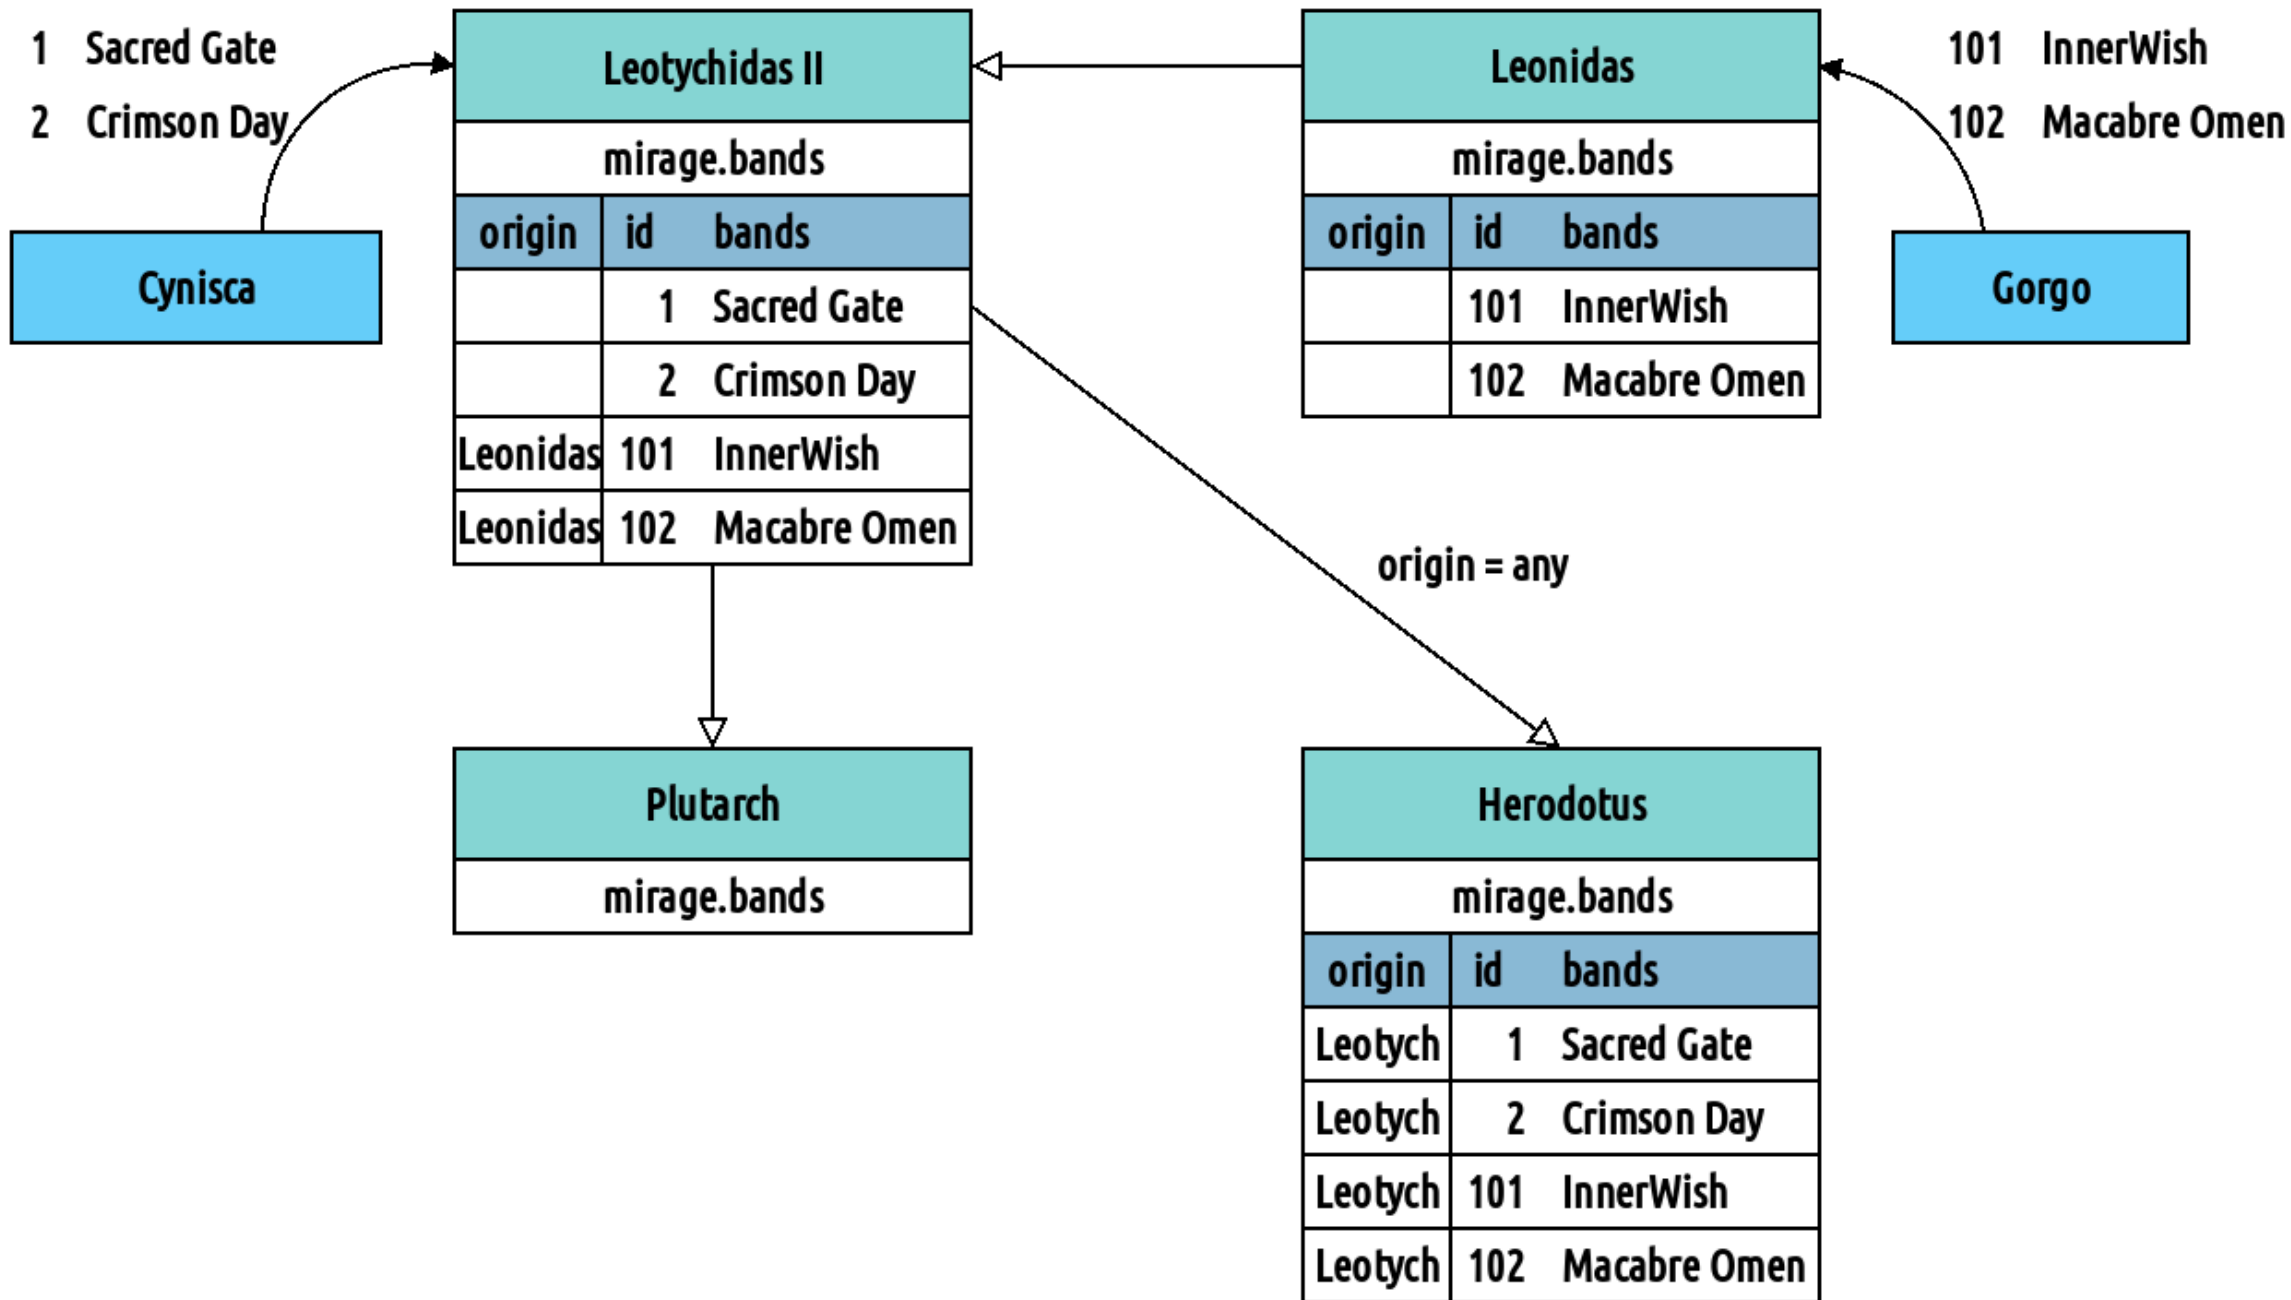


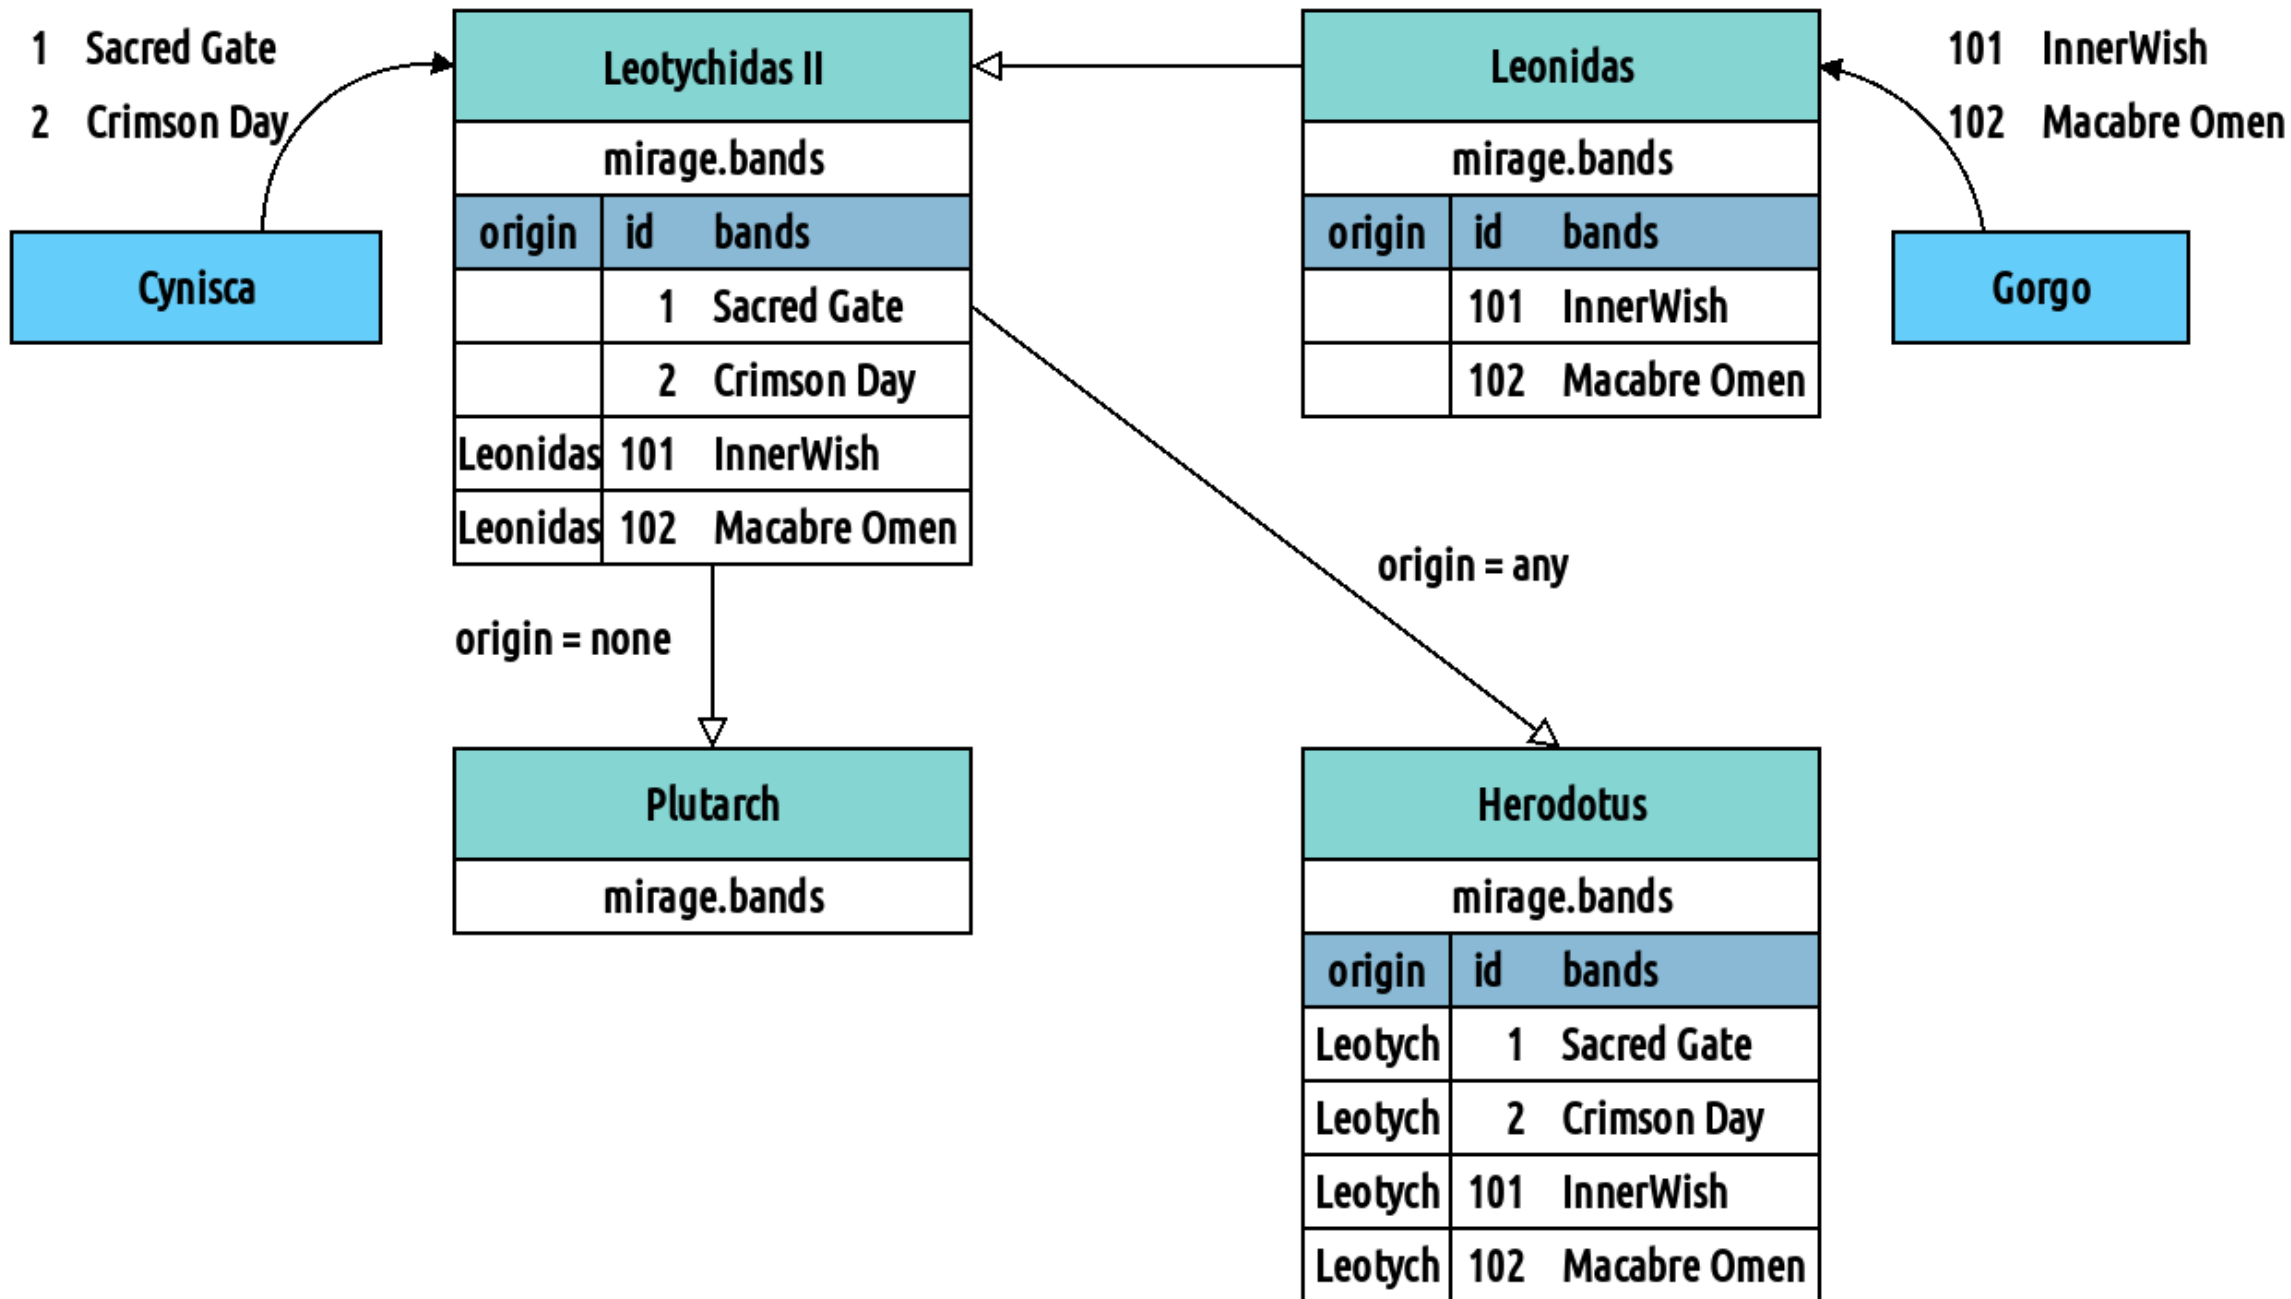


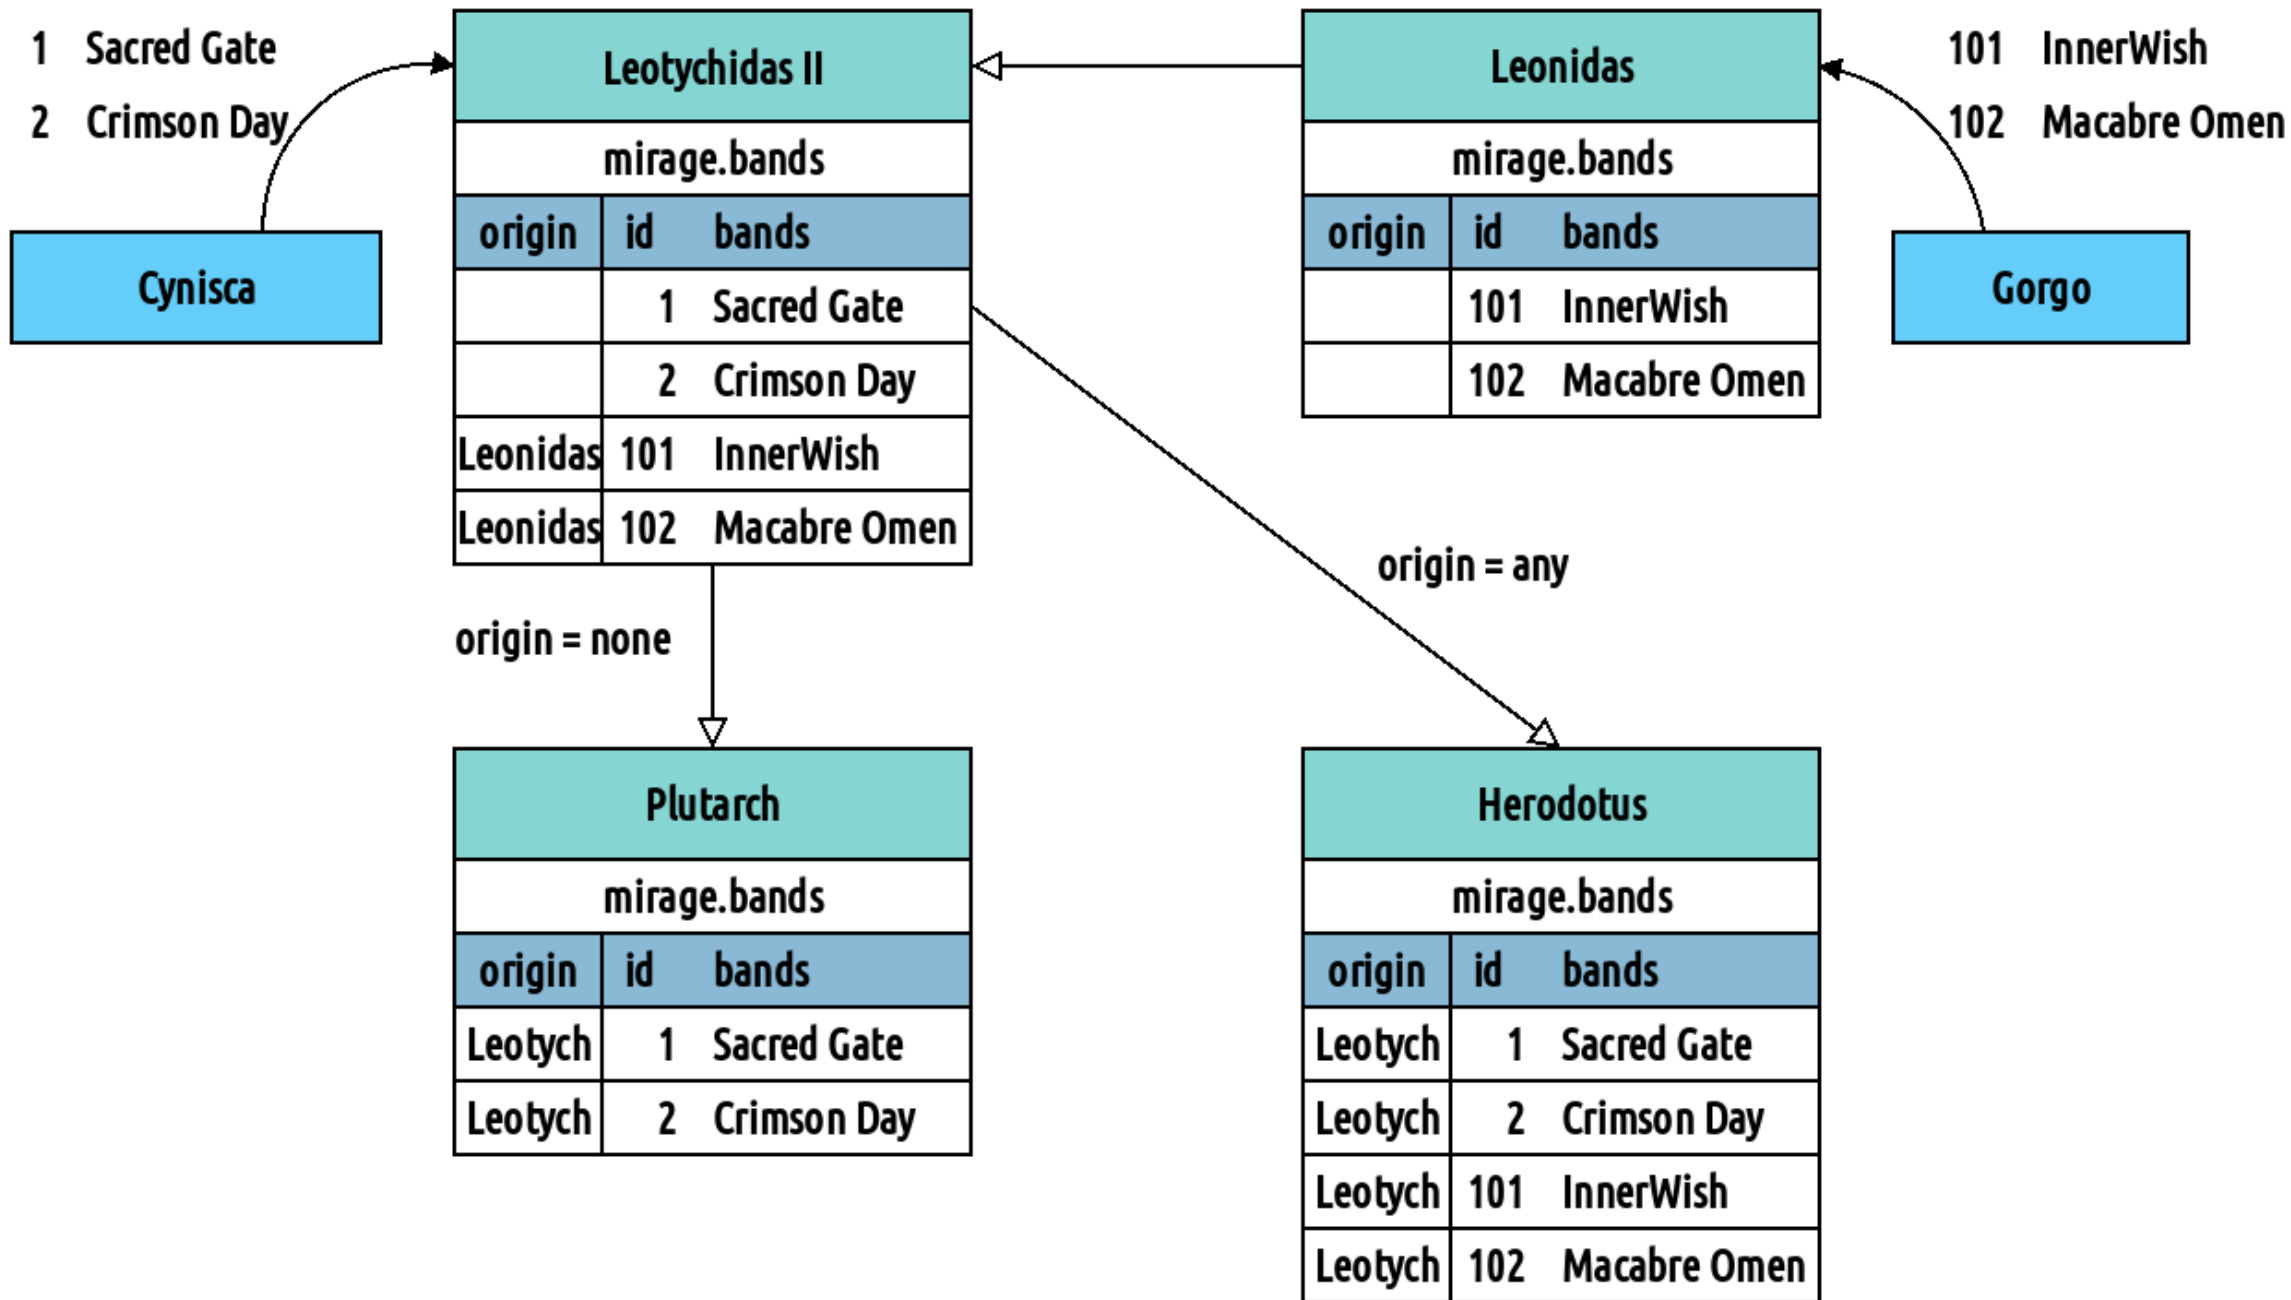


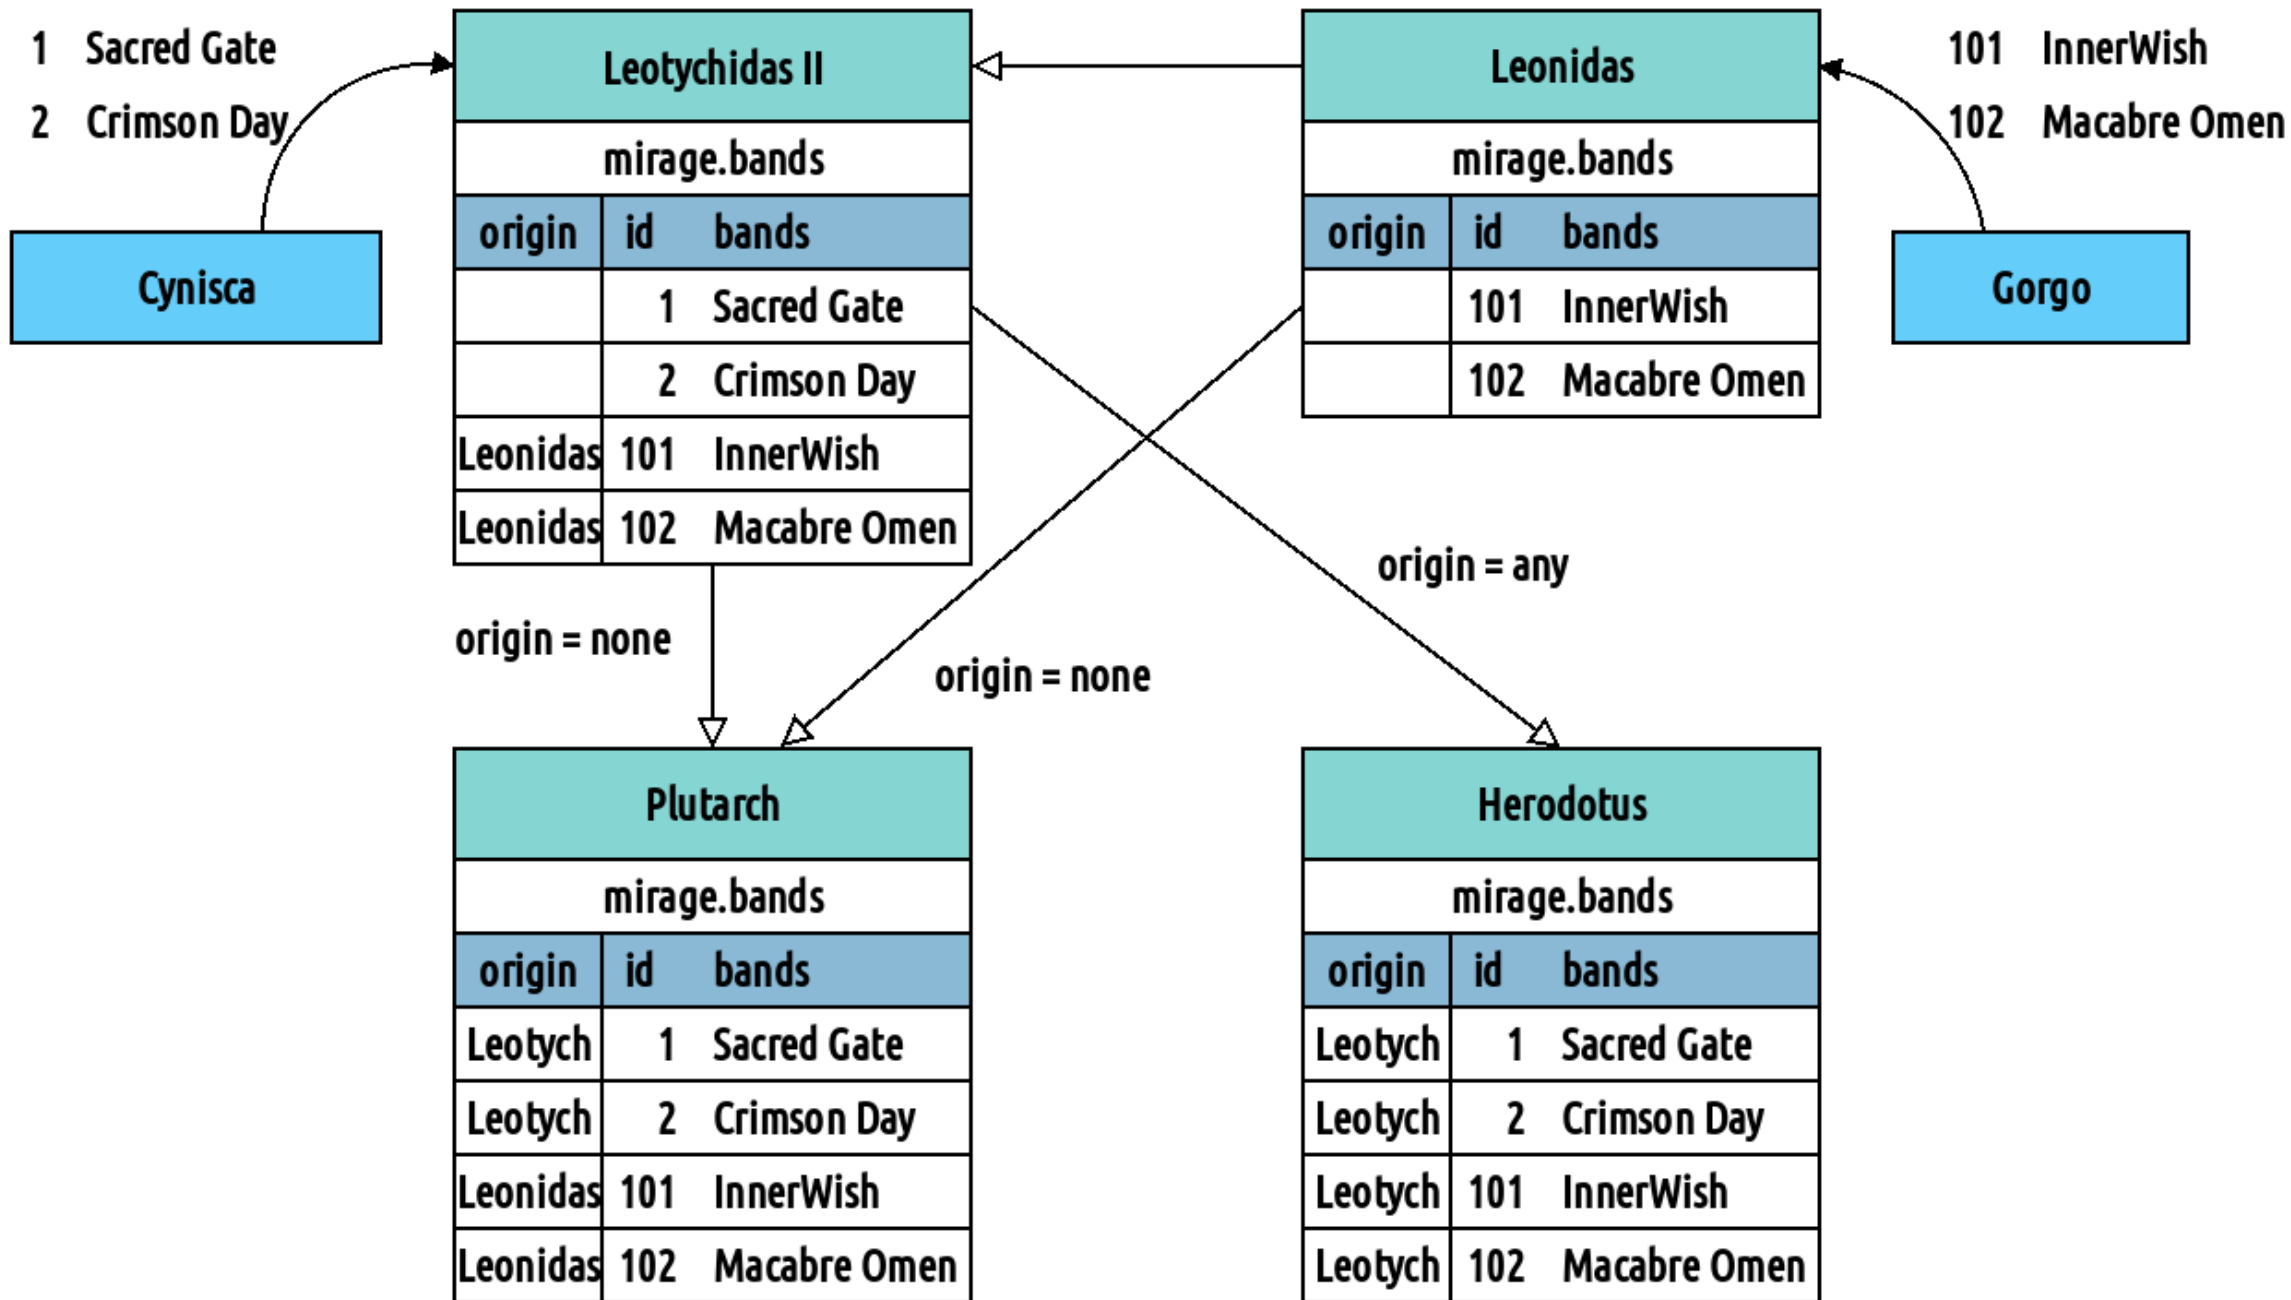











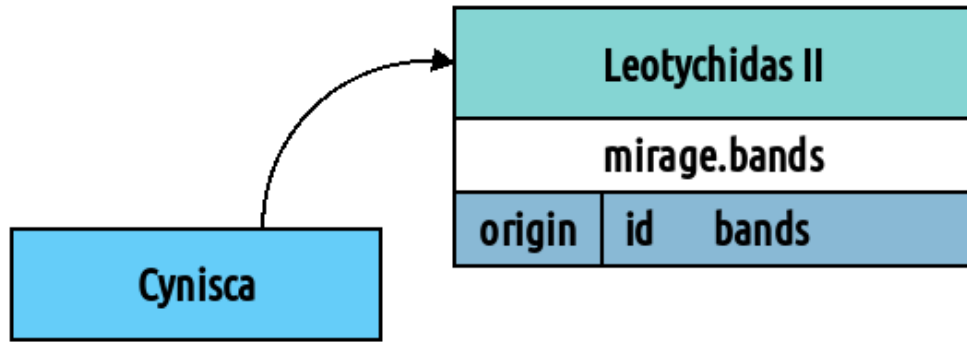



So that was
unidirectional logical
replication

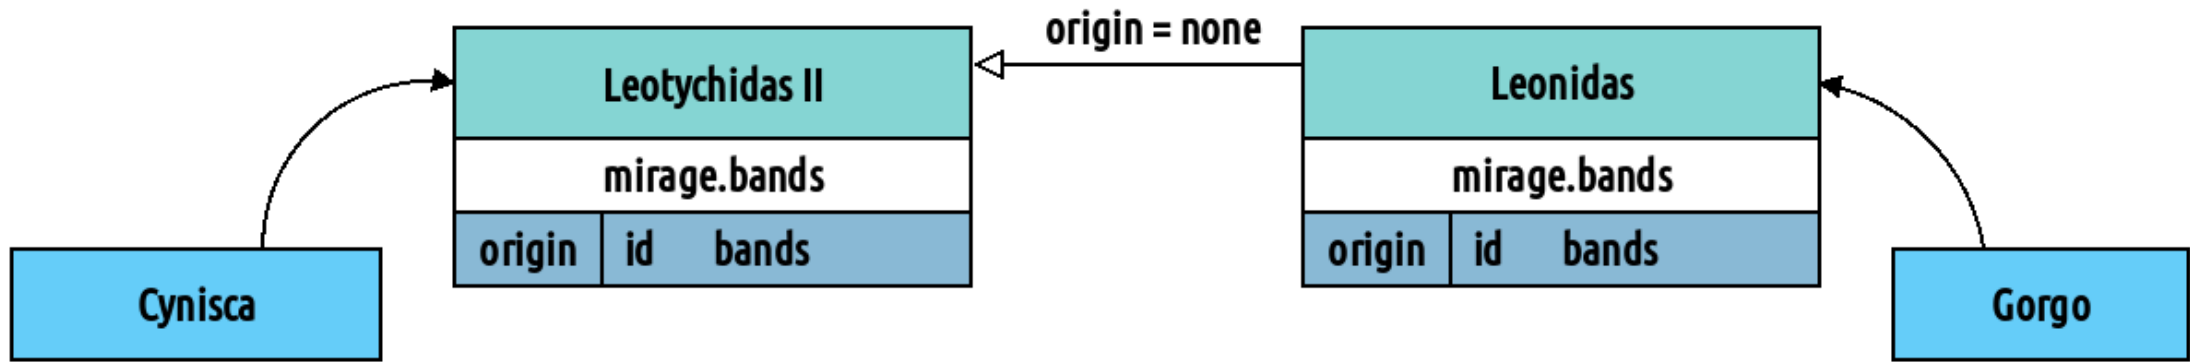


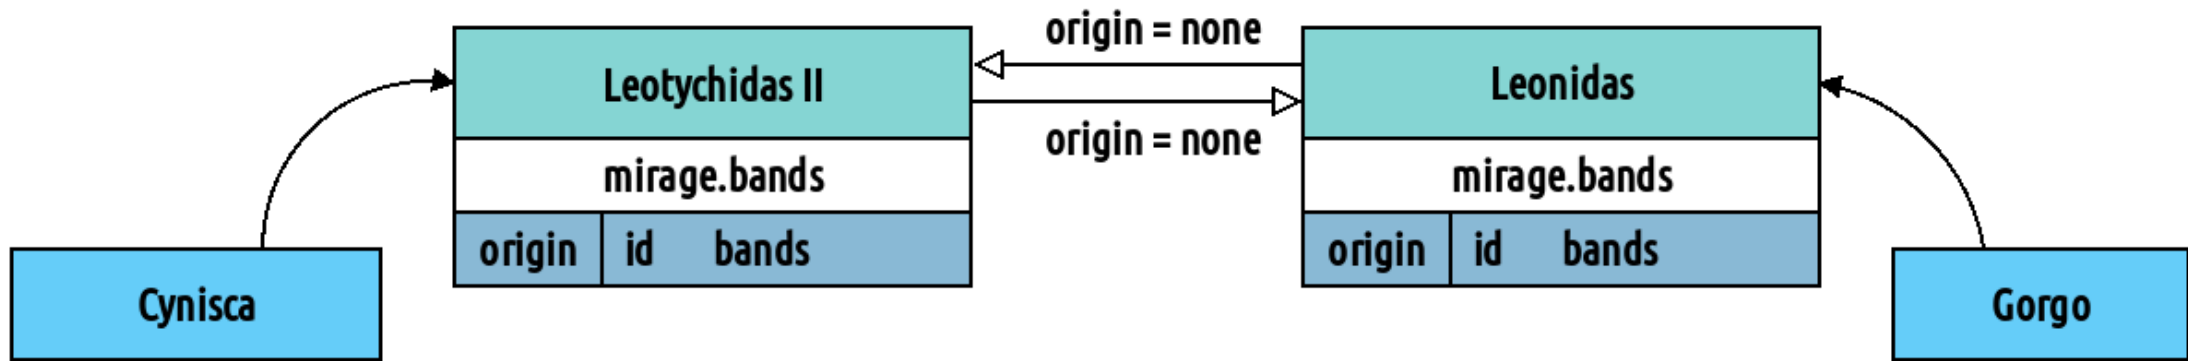
We are here for
active-active



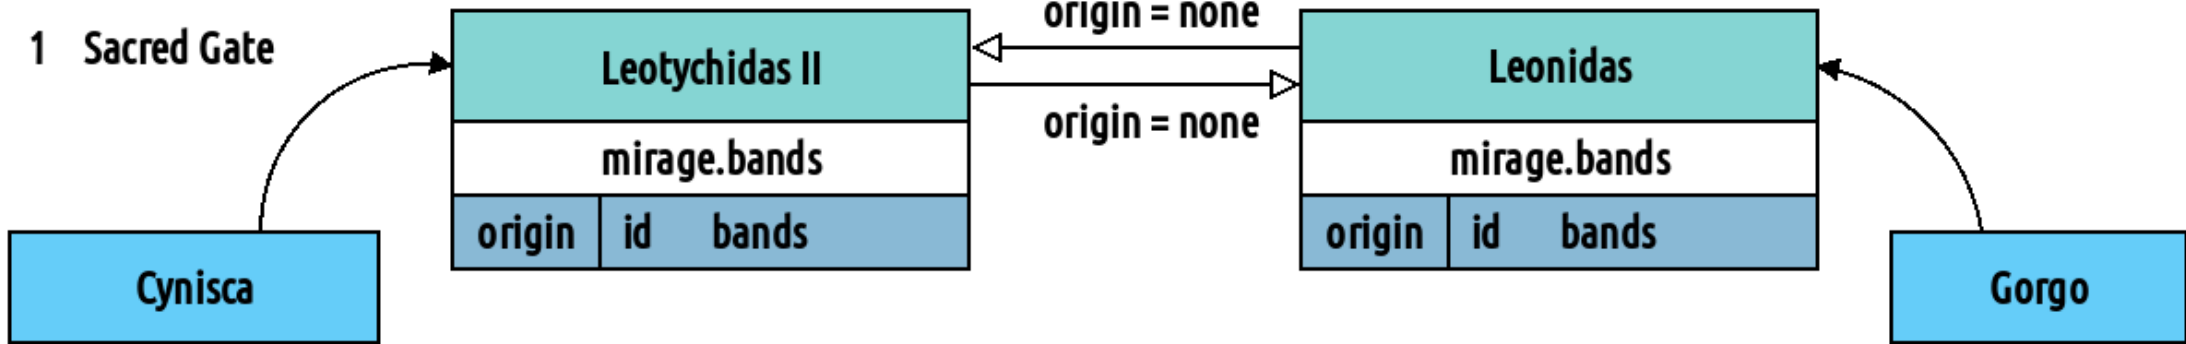


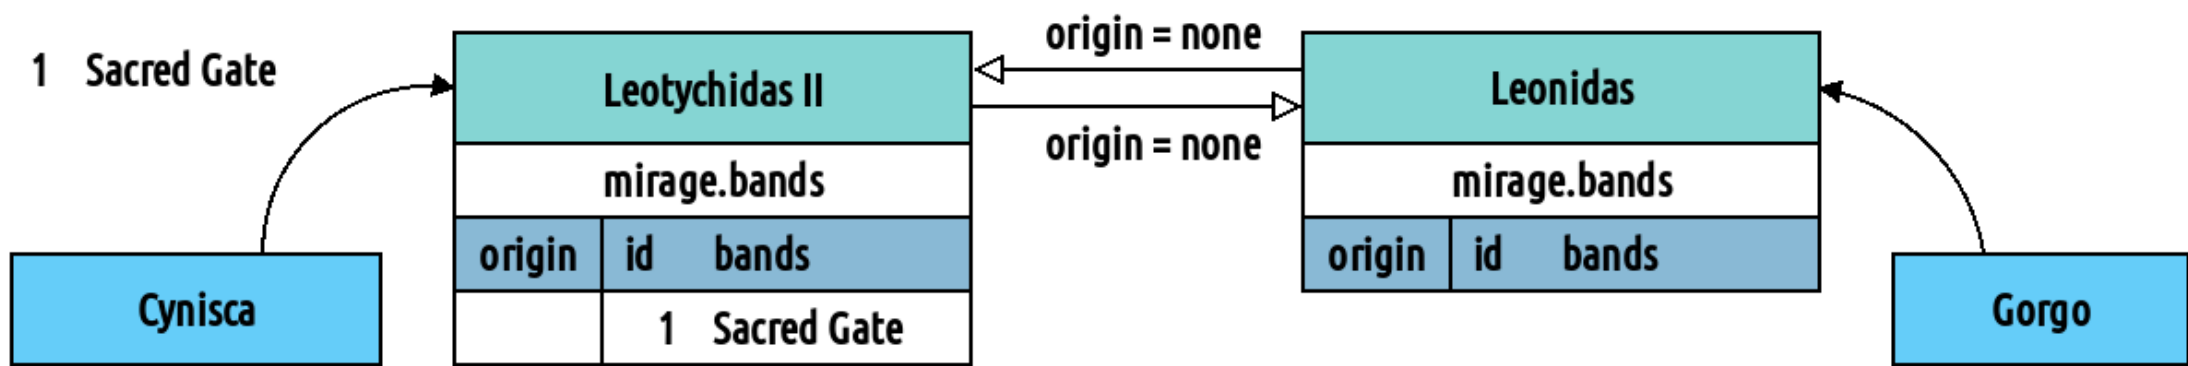


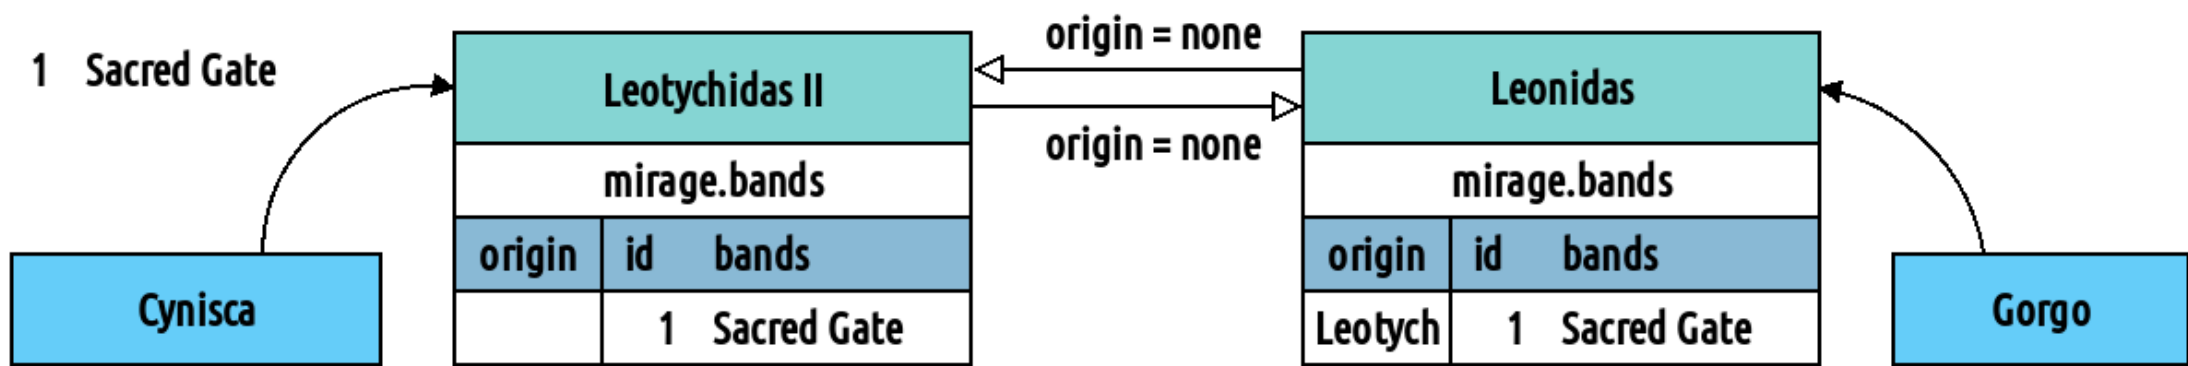


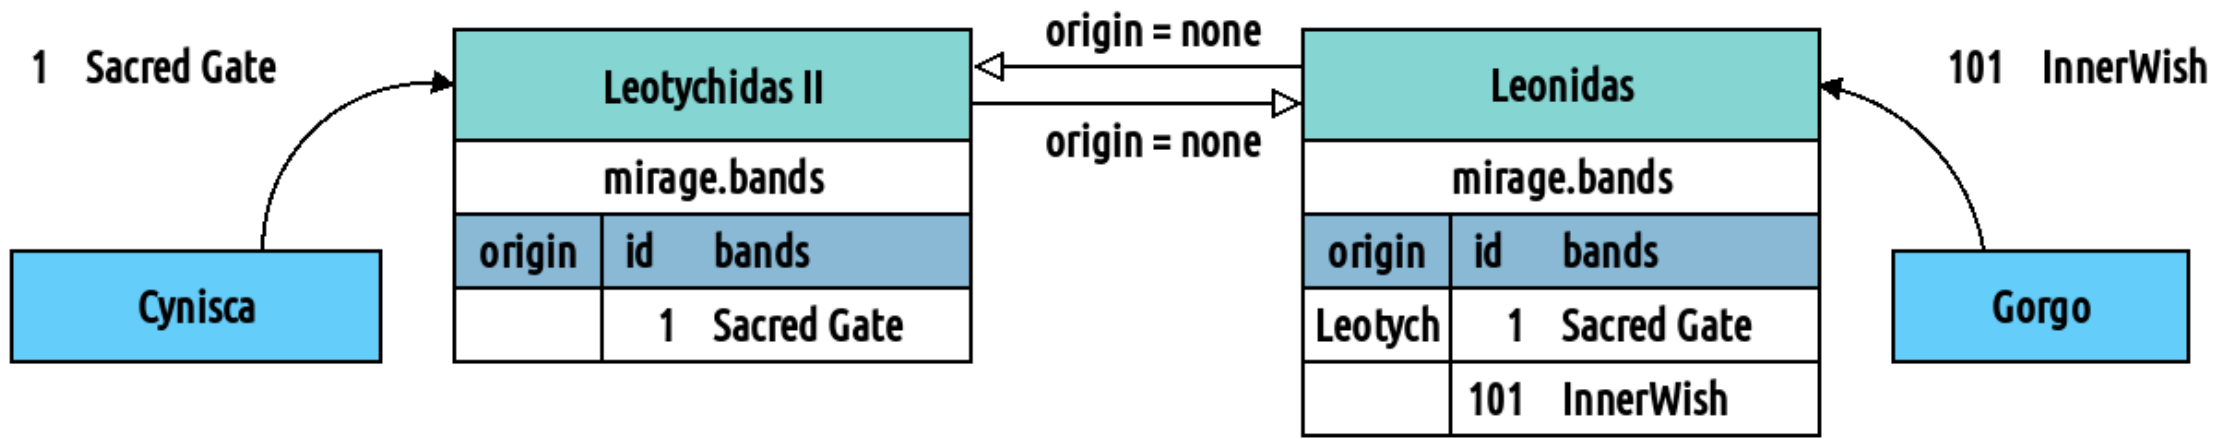


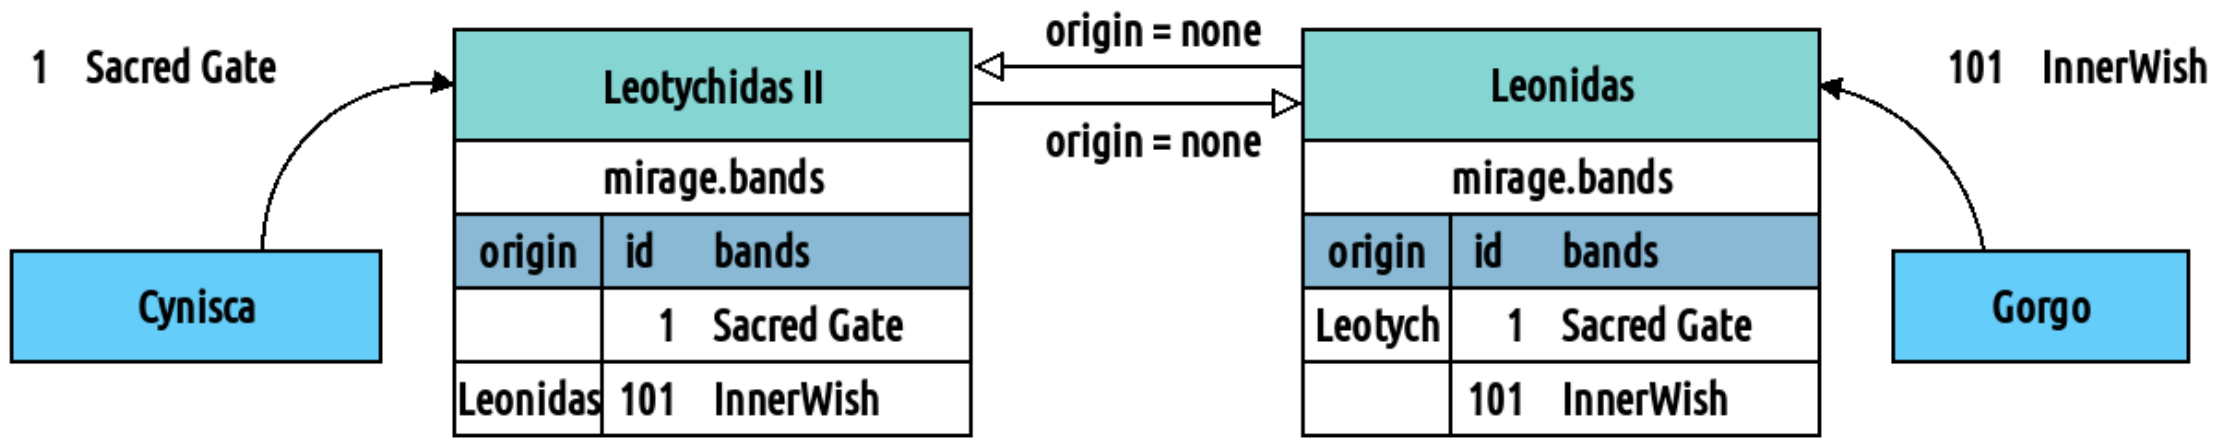
1 Sacred Gate

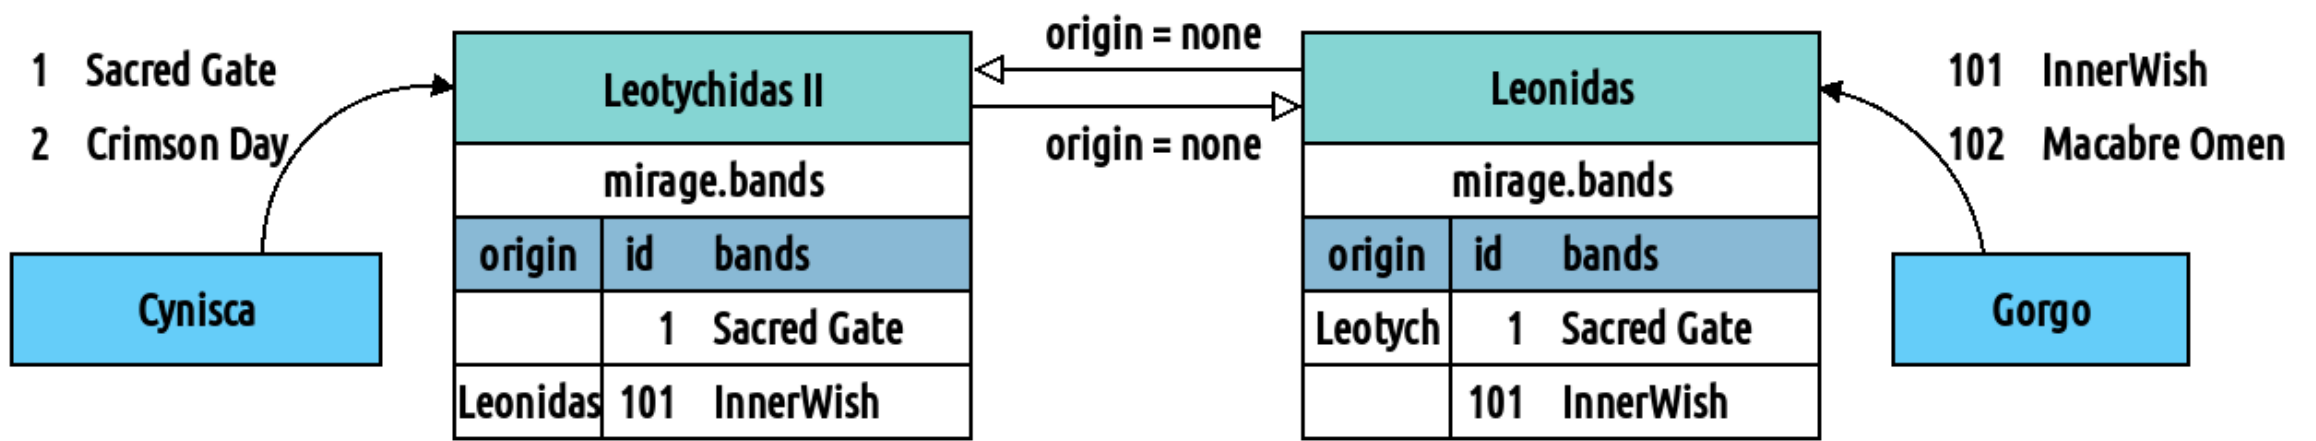


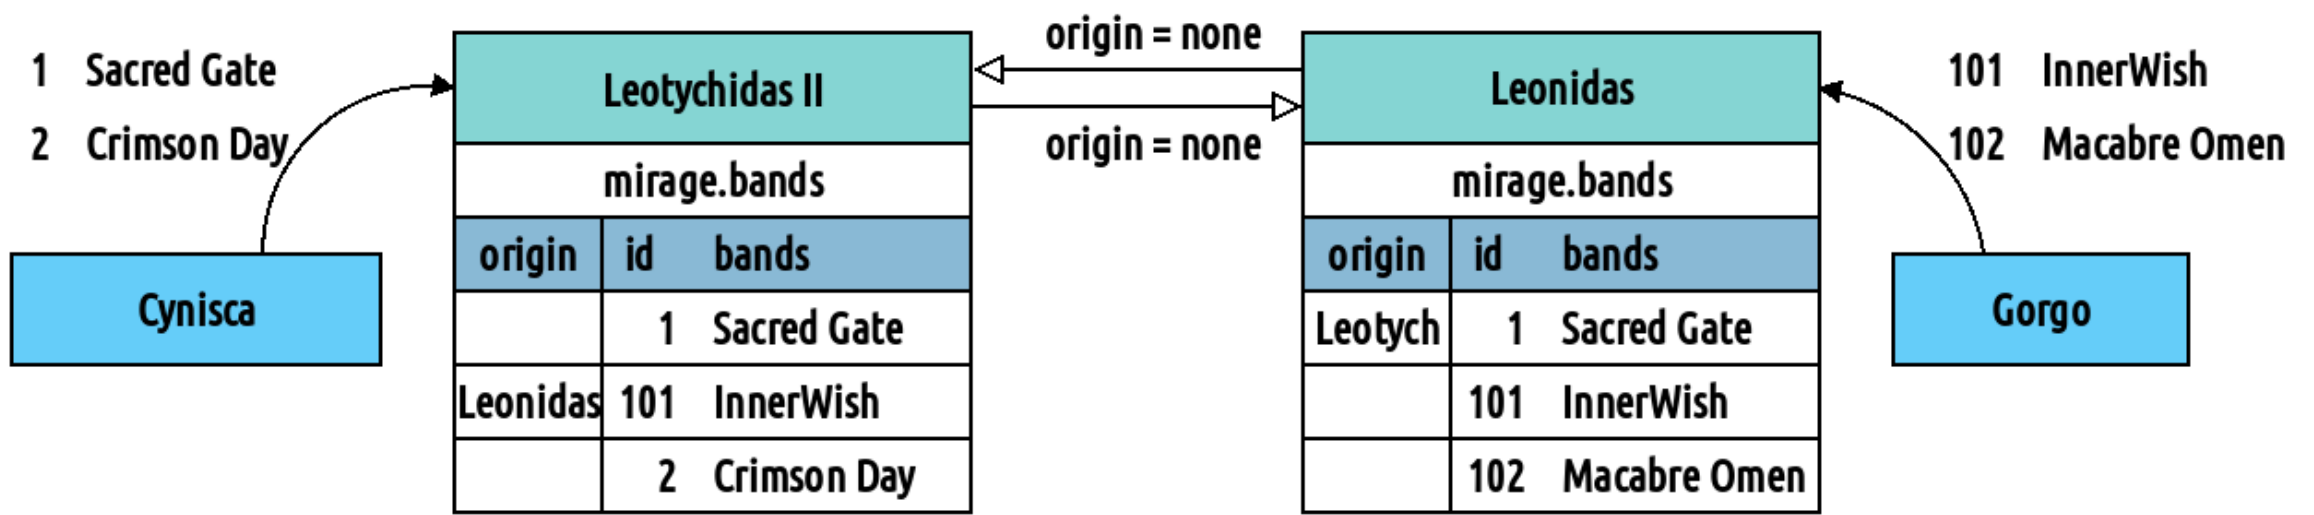















Conflicts



Conflicts between Kings of Sparta


Demaratus and Cleomenes I (late 6th early 5th cent BCE)


Agesilaus and Leotychidas II (4th century BCE)


Cleombrotus II and Leonidas II (3rd century BCE)

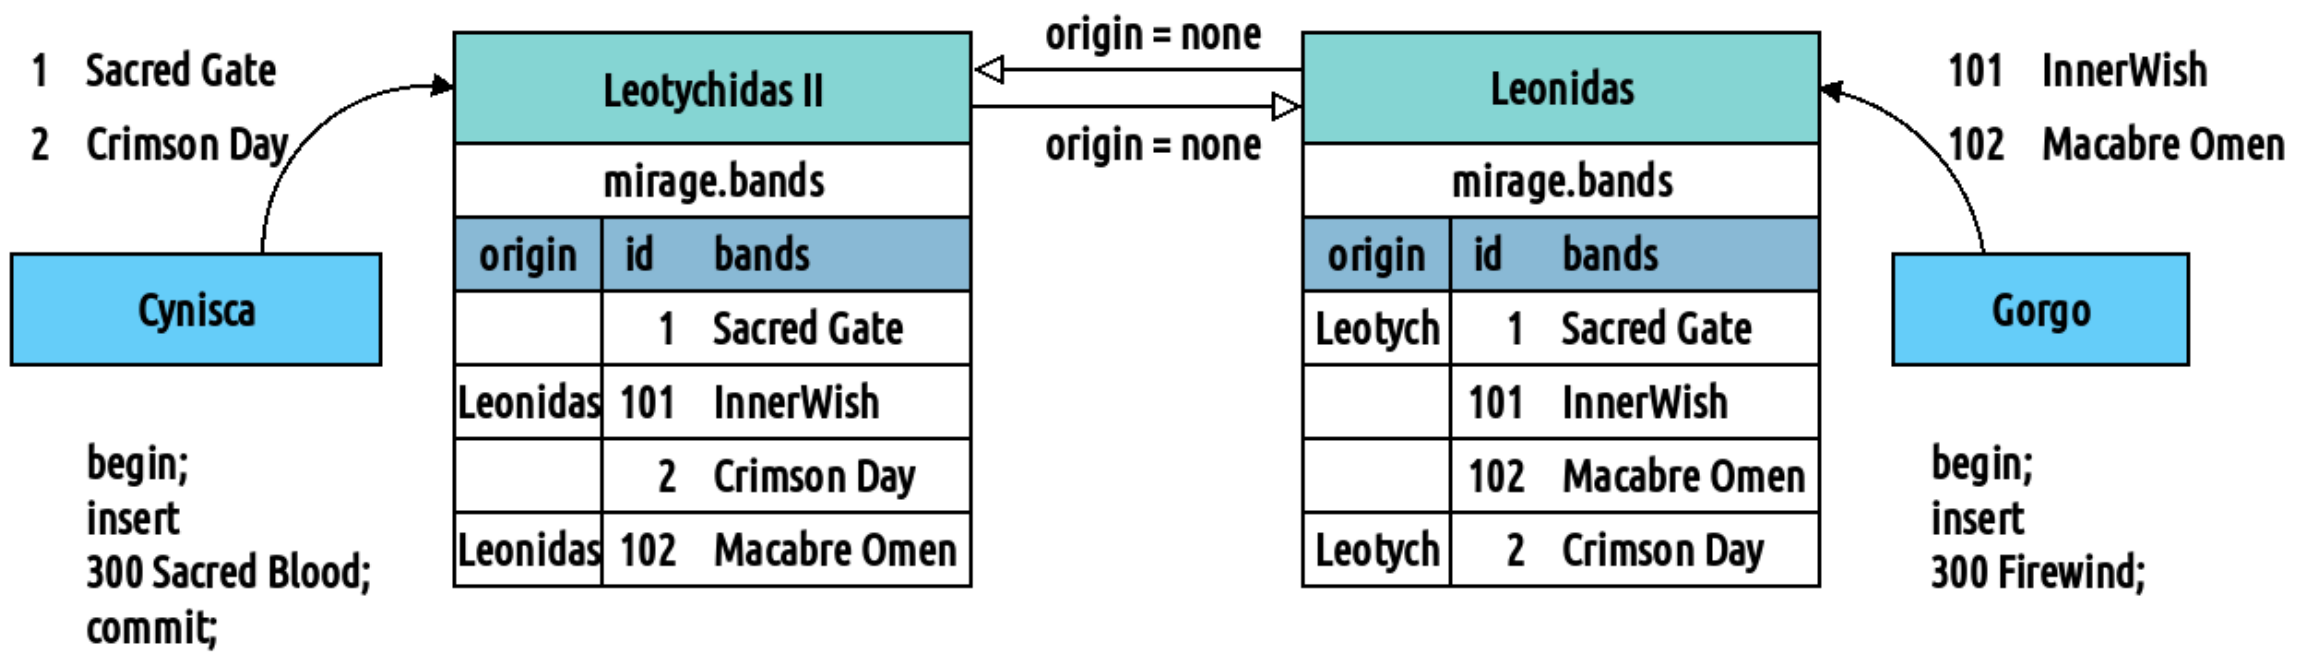
Archidamus V and Cleomenes III (late 3rd cent BCE)

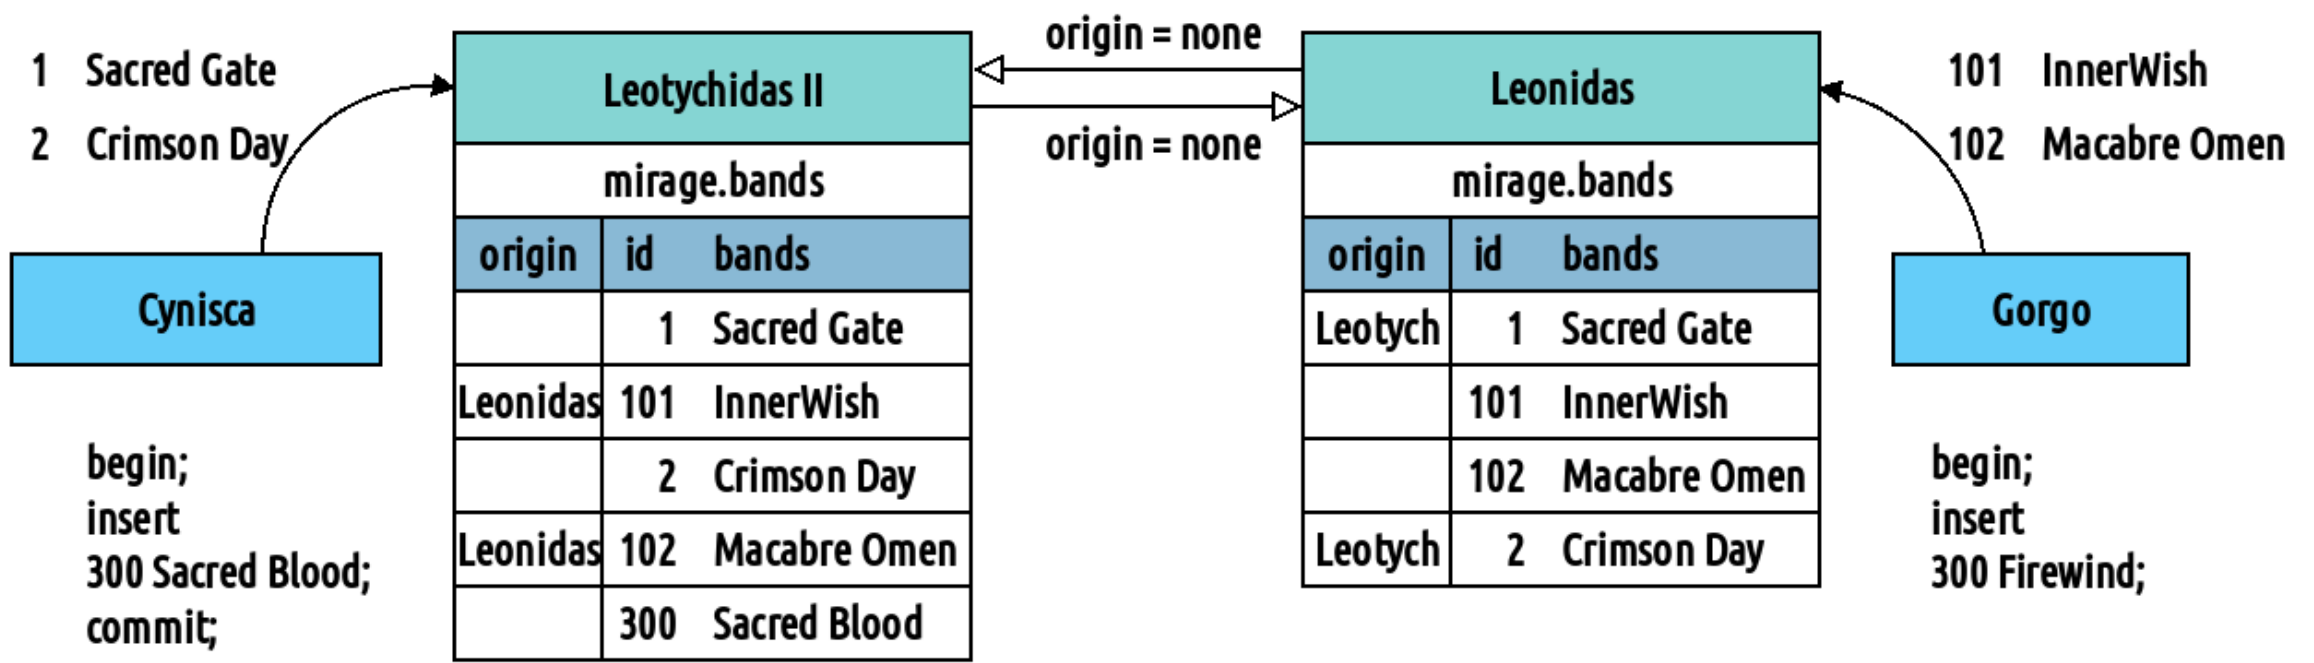


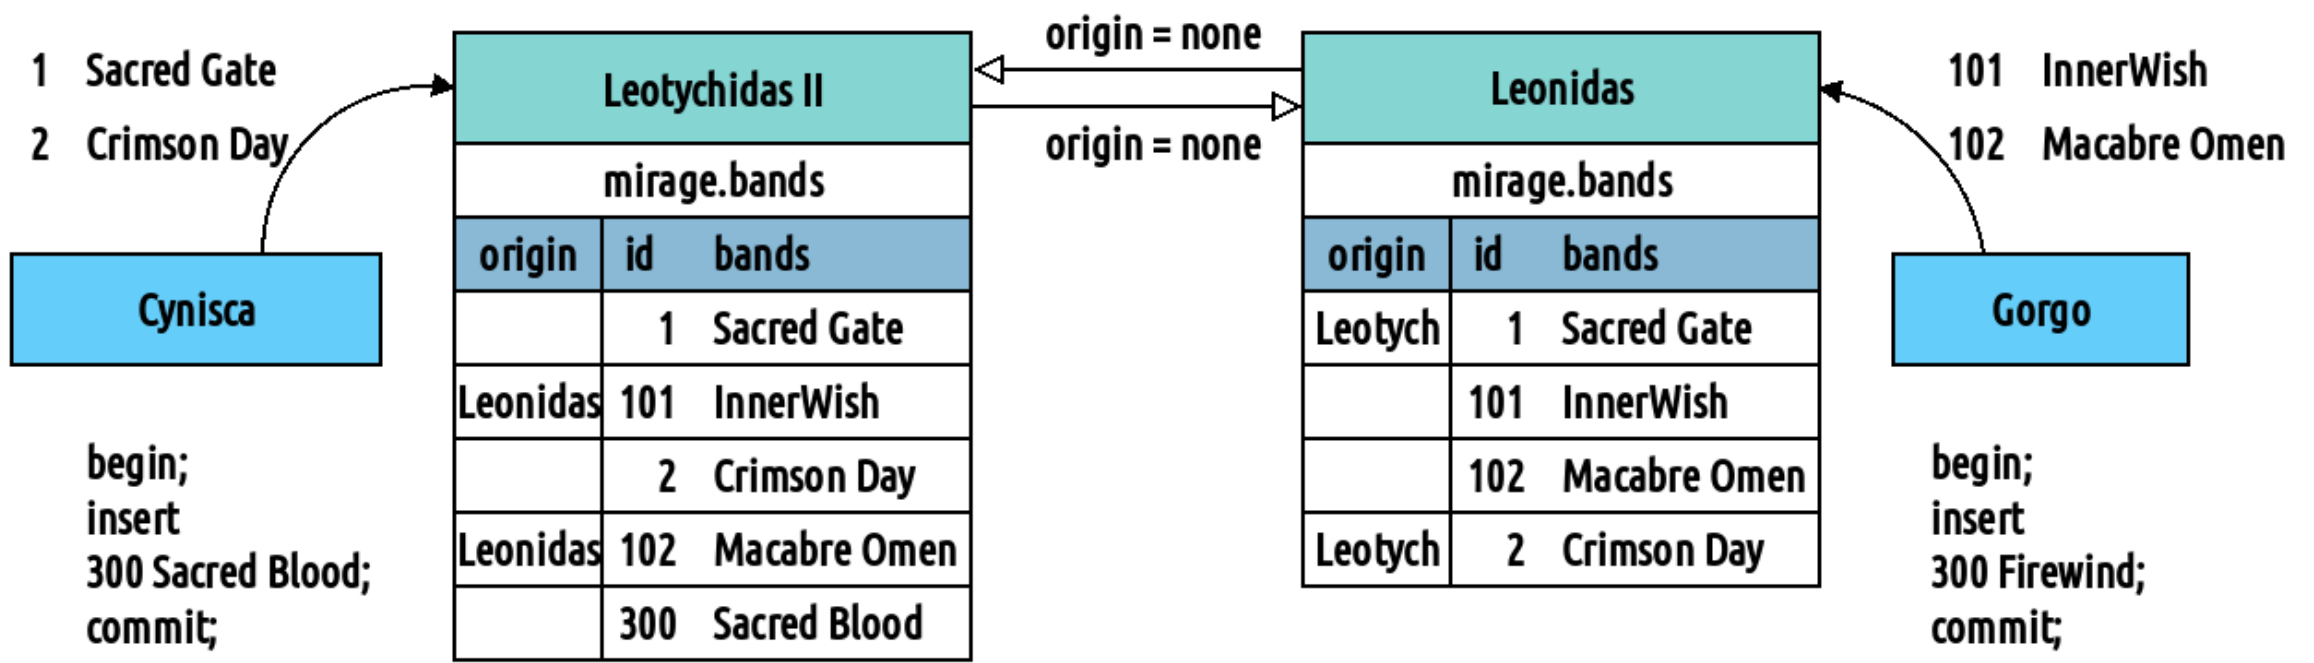


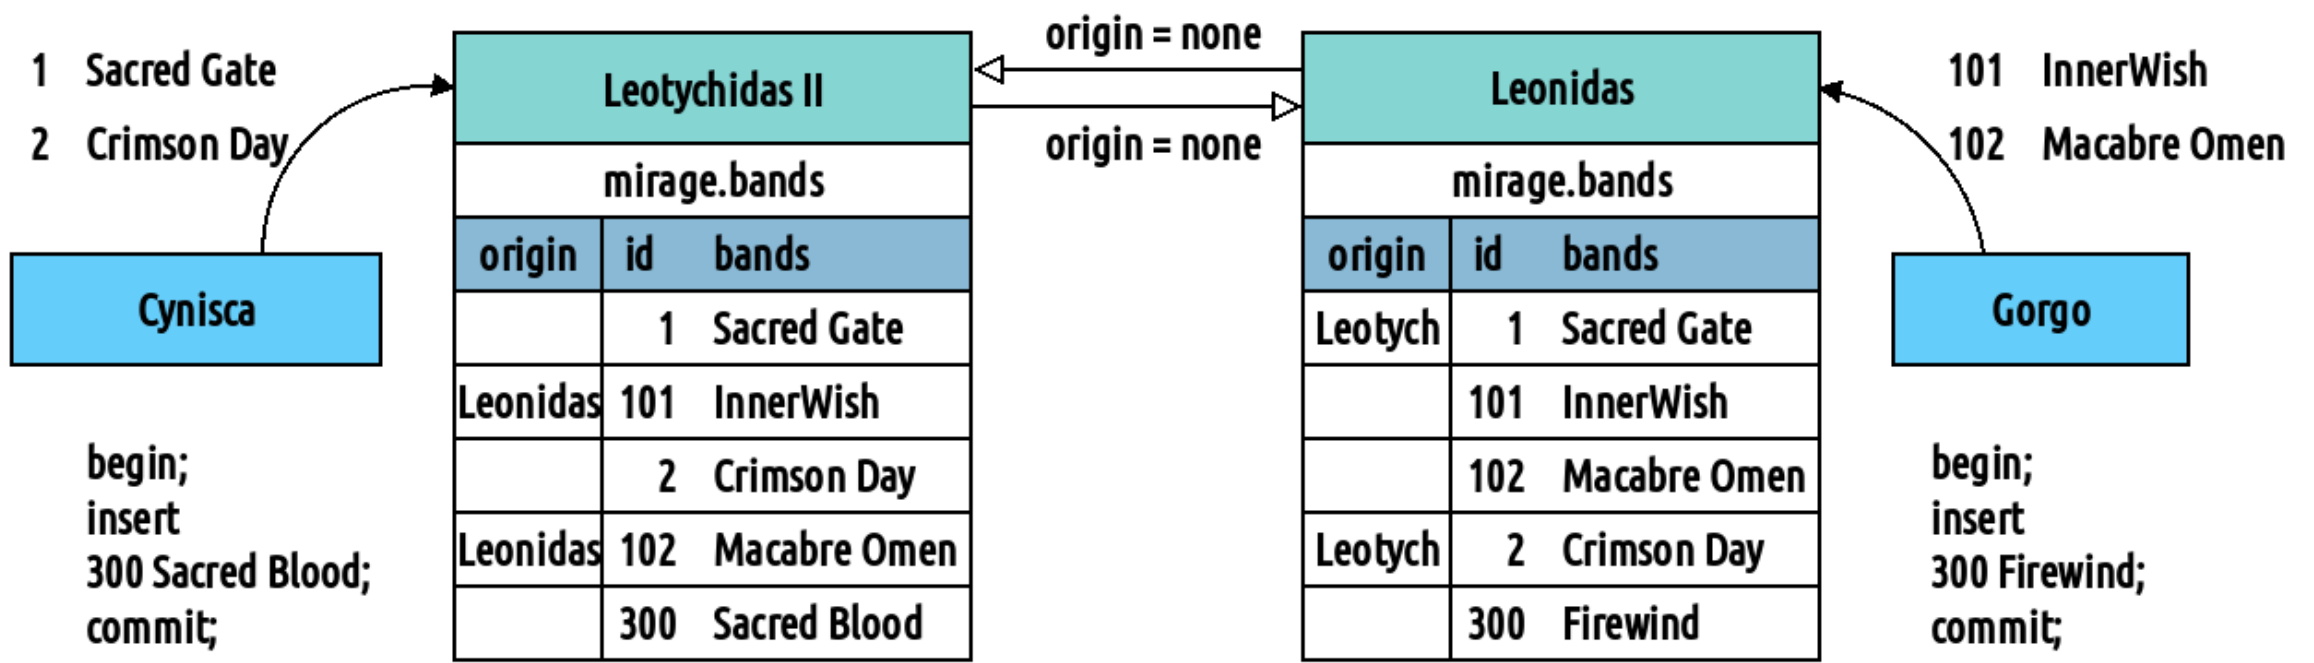


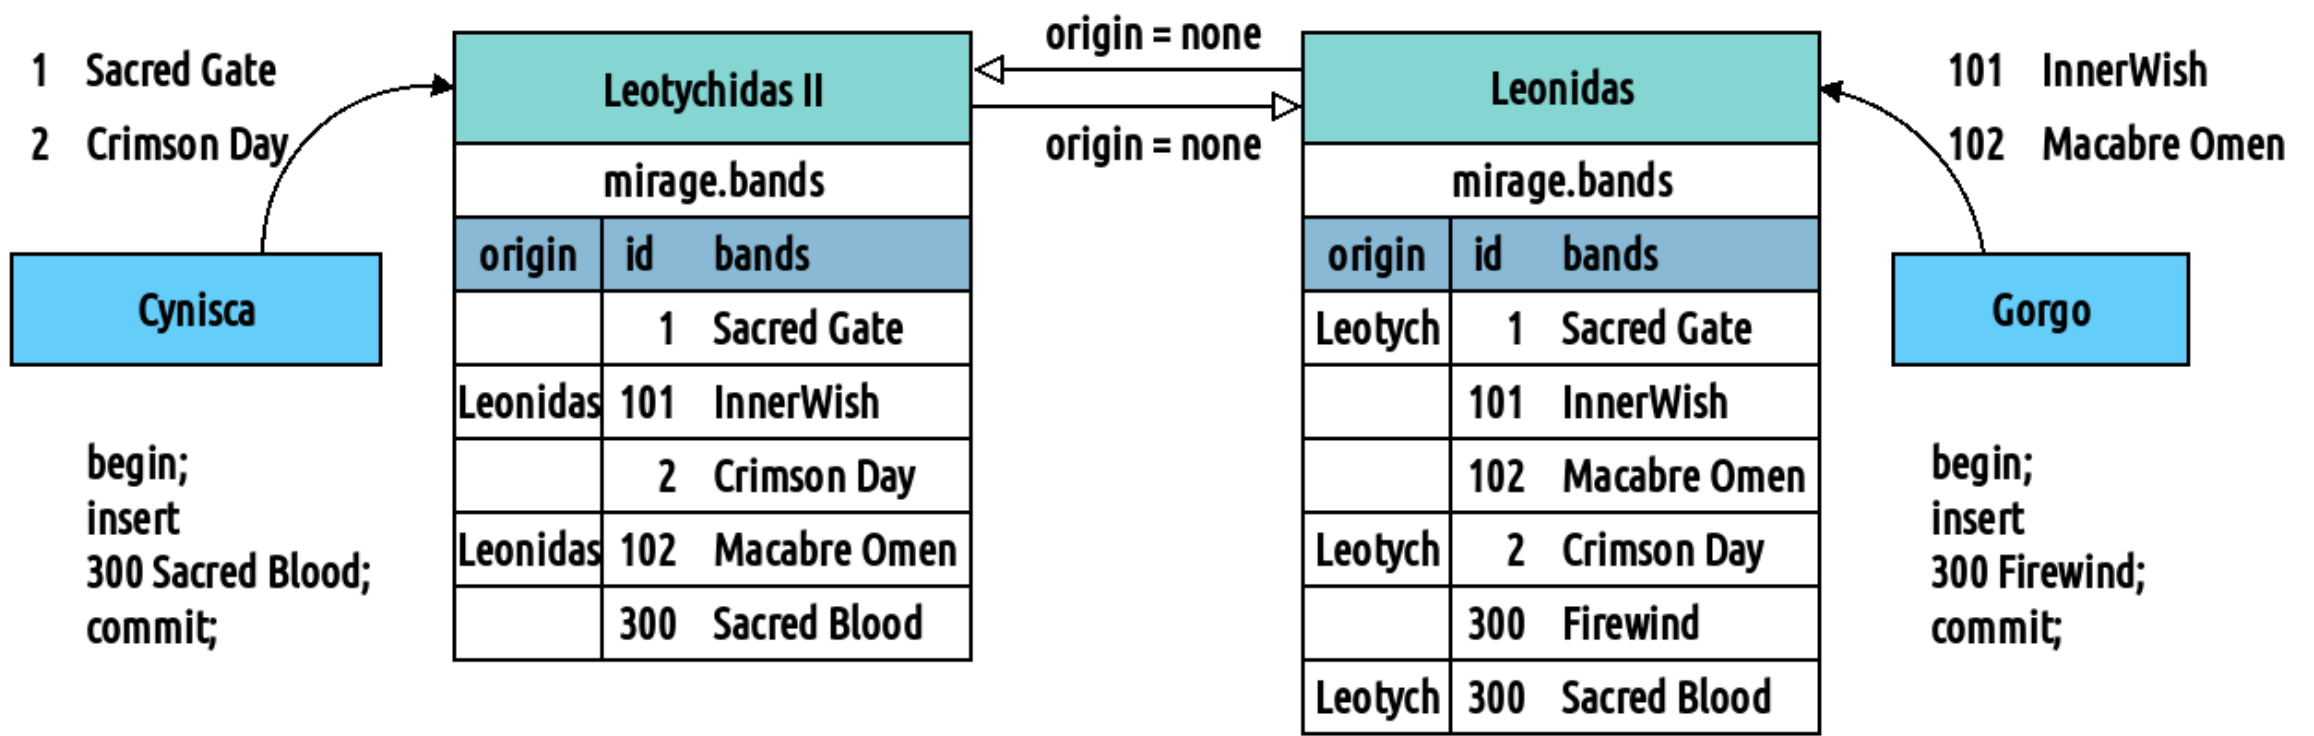


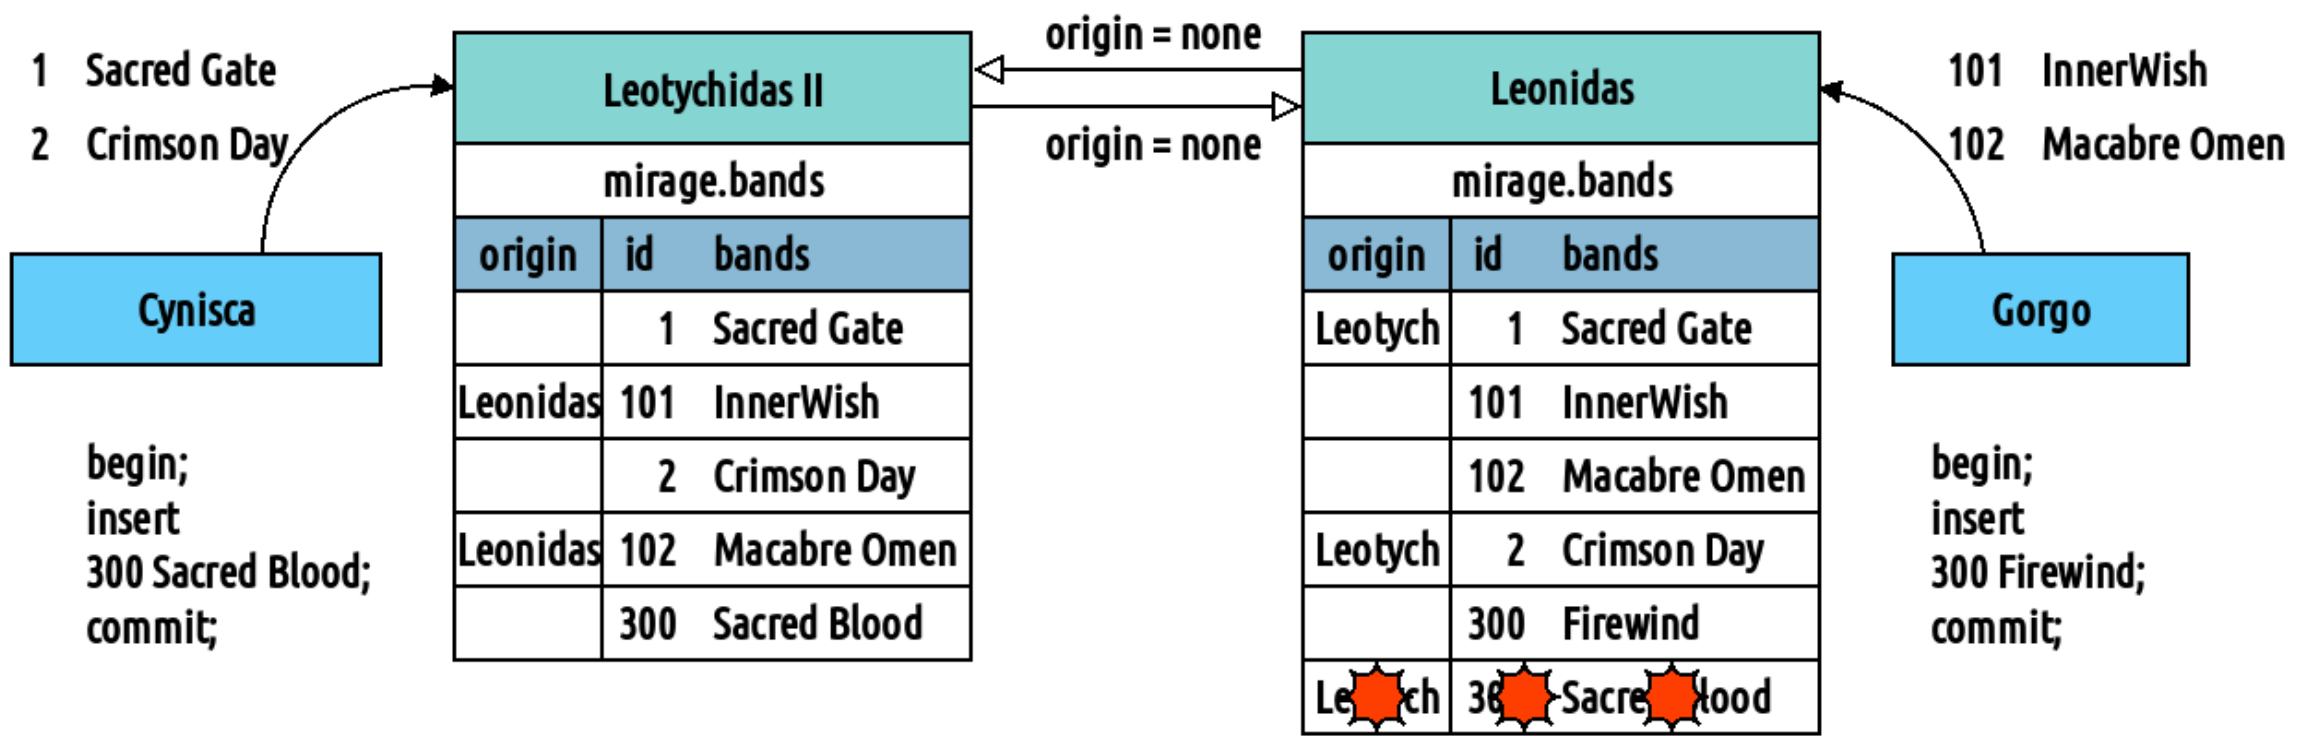


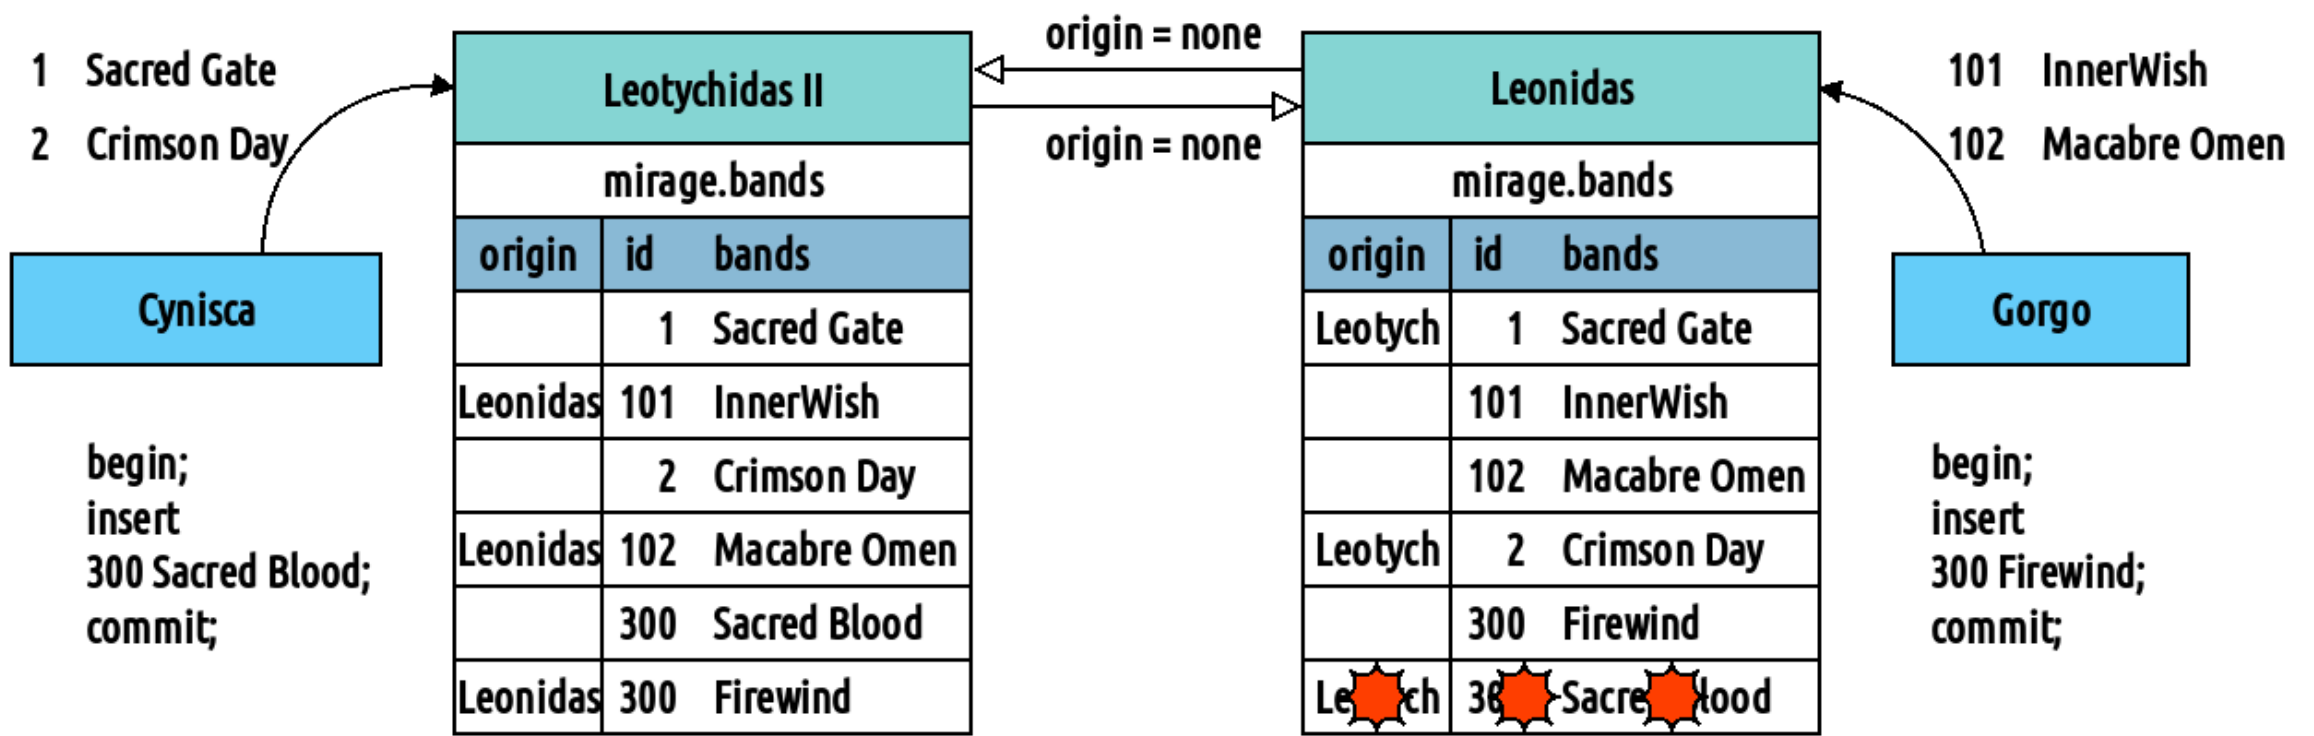


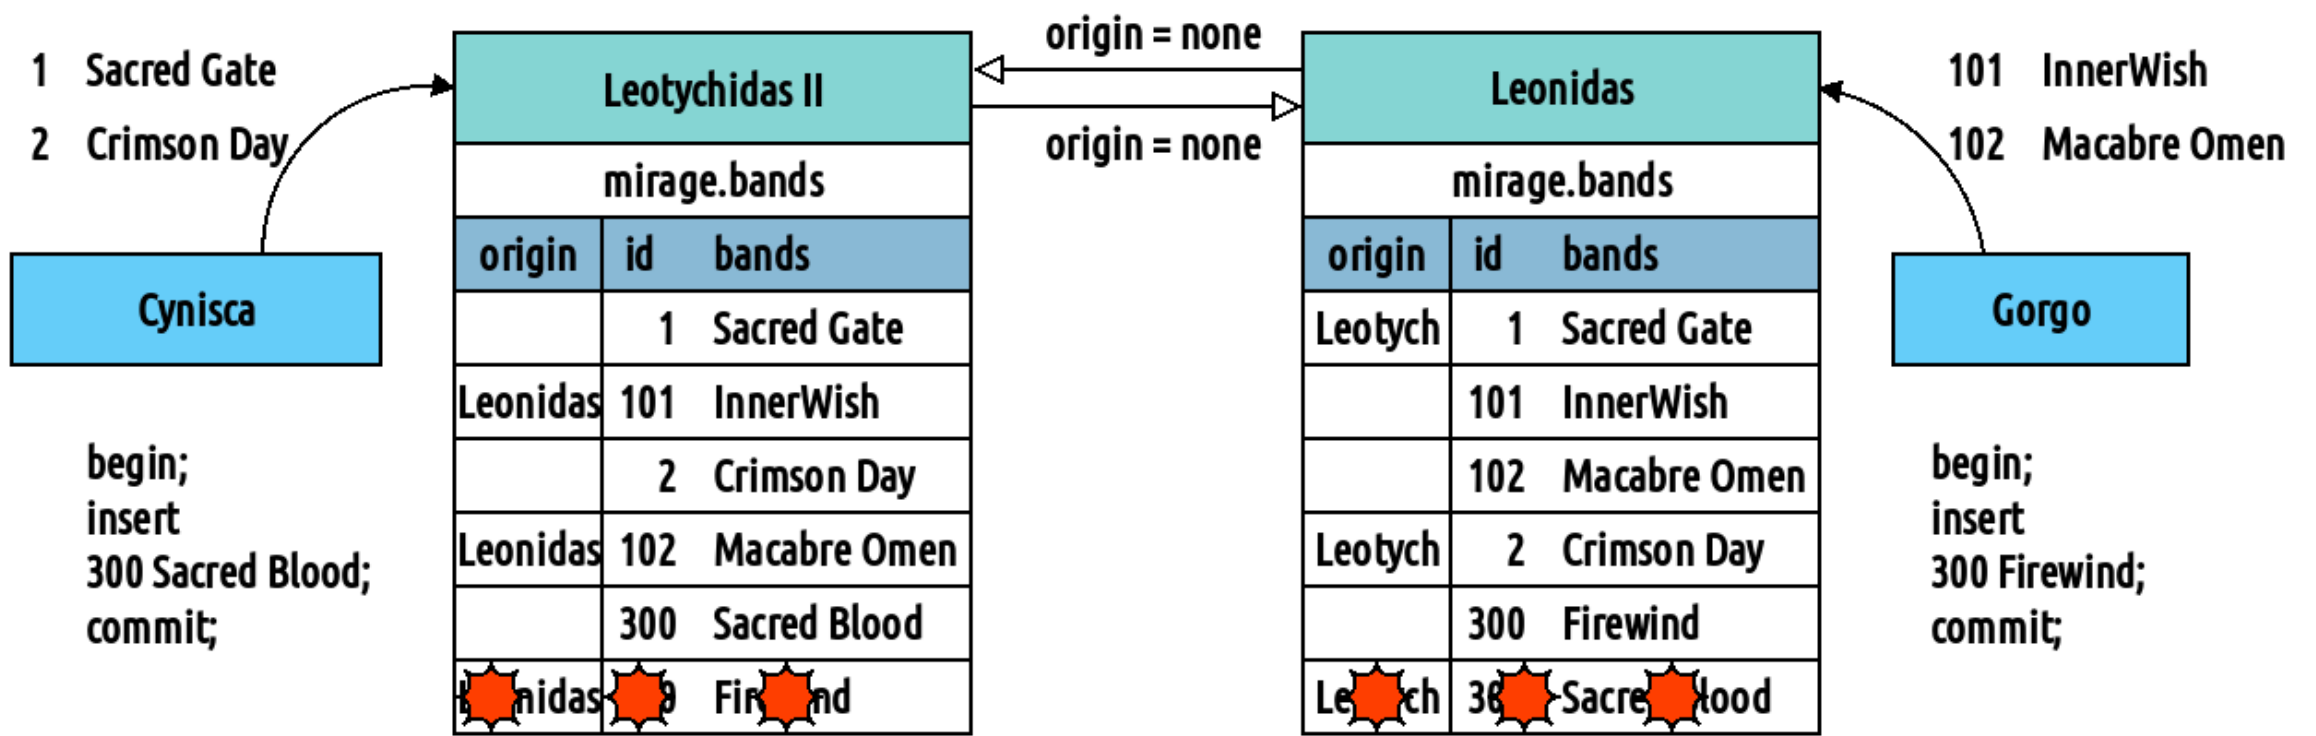


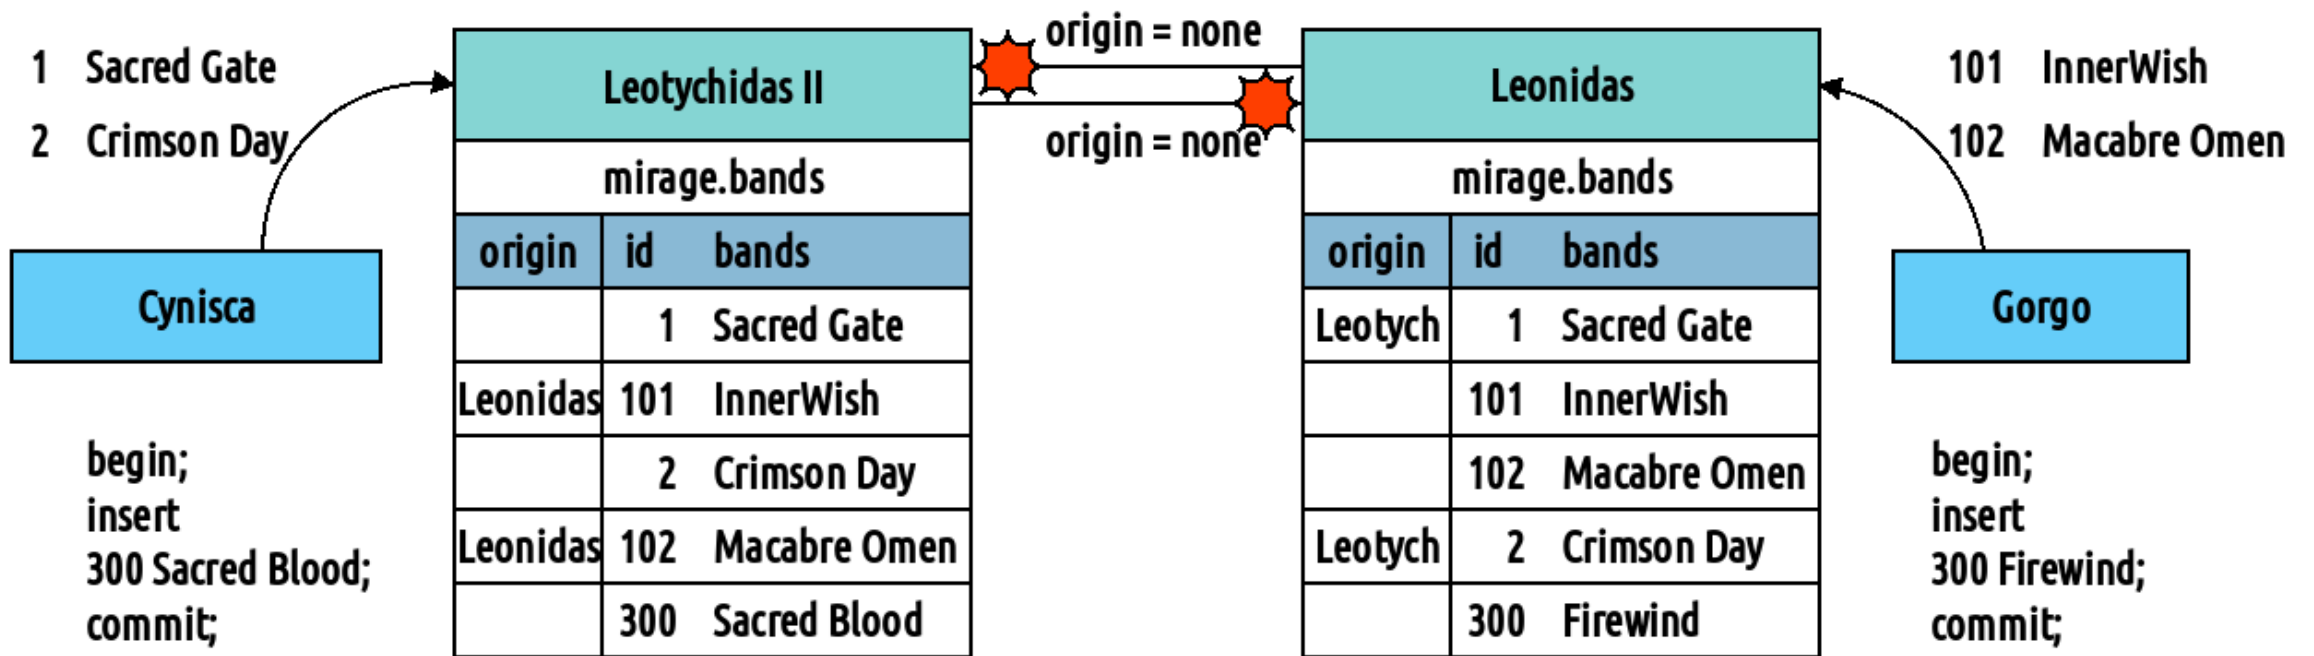


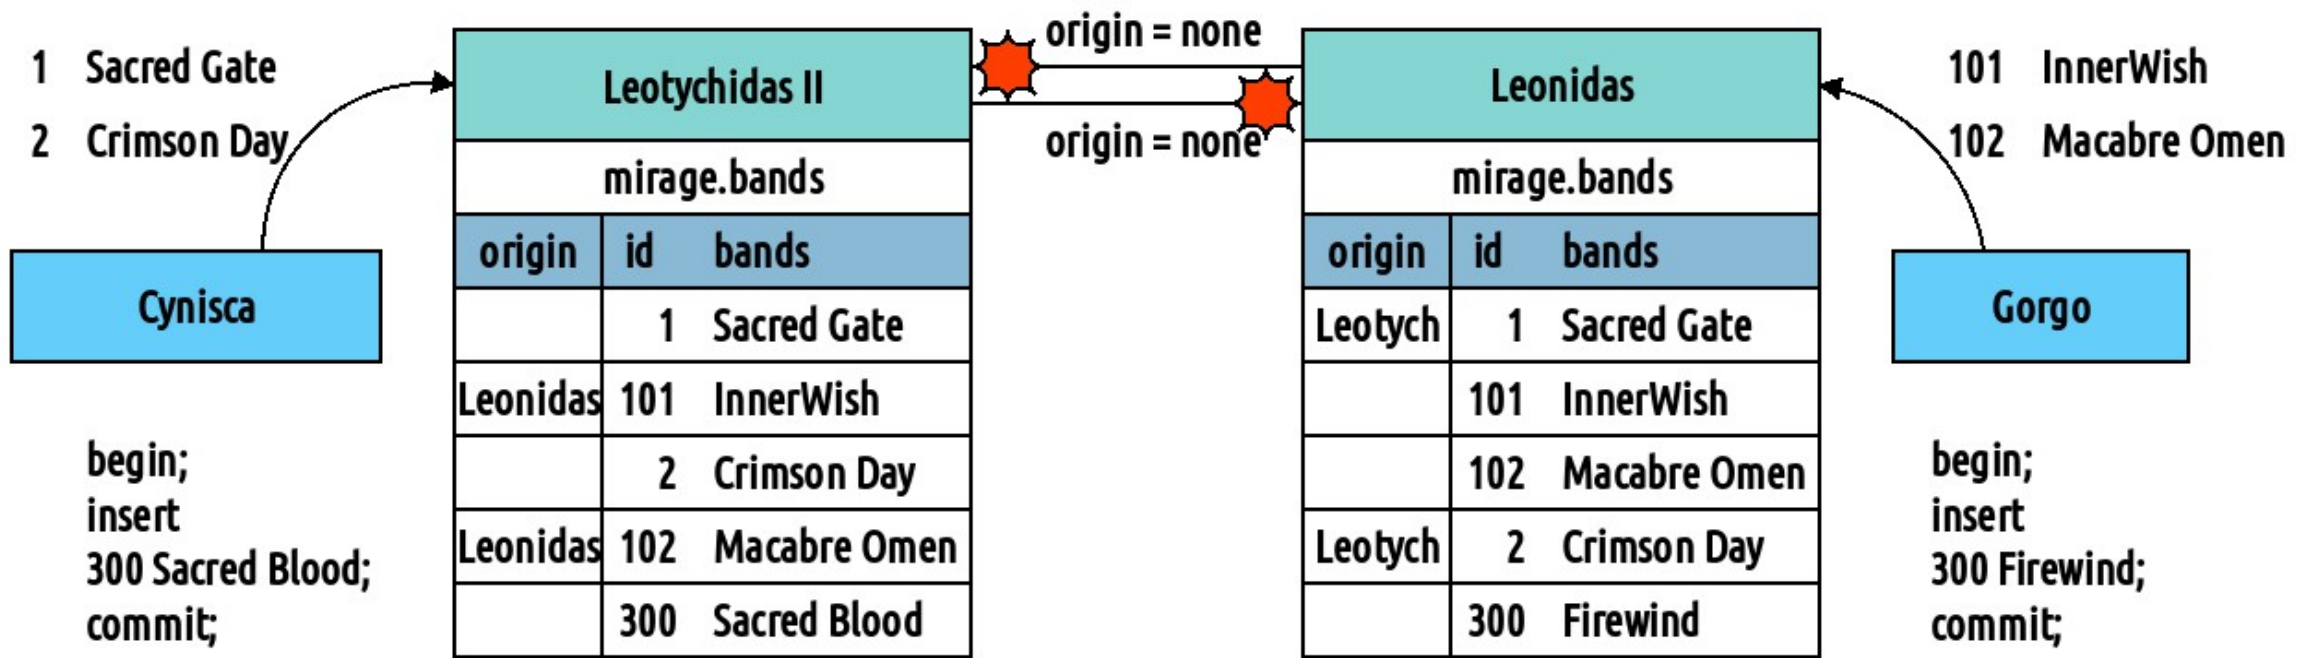












A C I D

No Spartan King (regent) is above the law



No Spartan King (regent) is above the law

Pausanias, regent for Pleistarchus (479-478 BCE)

Allegations for collaboration with Persia




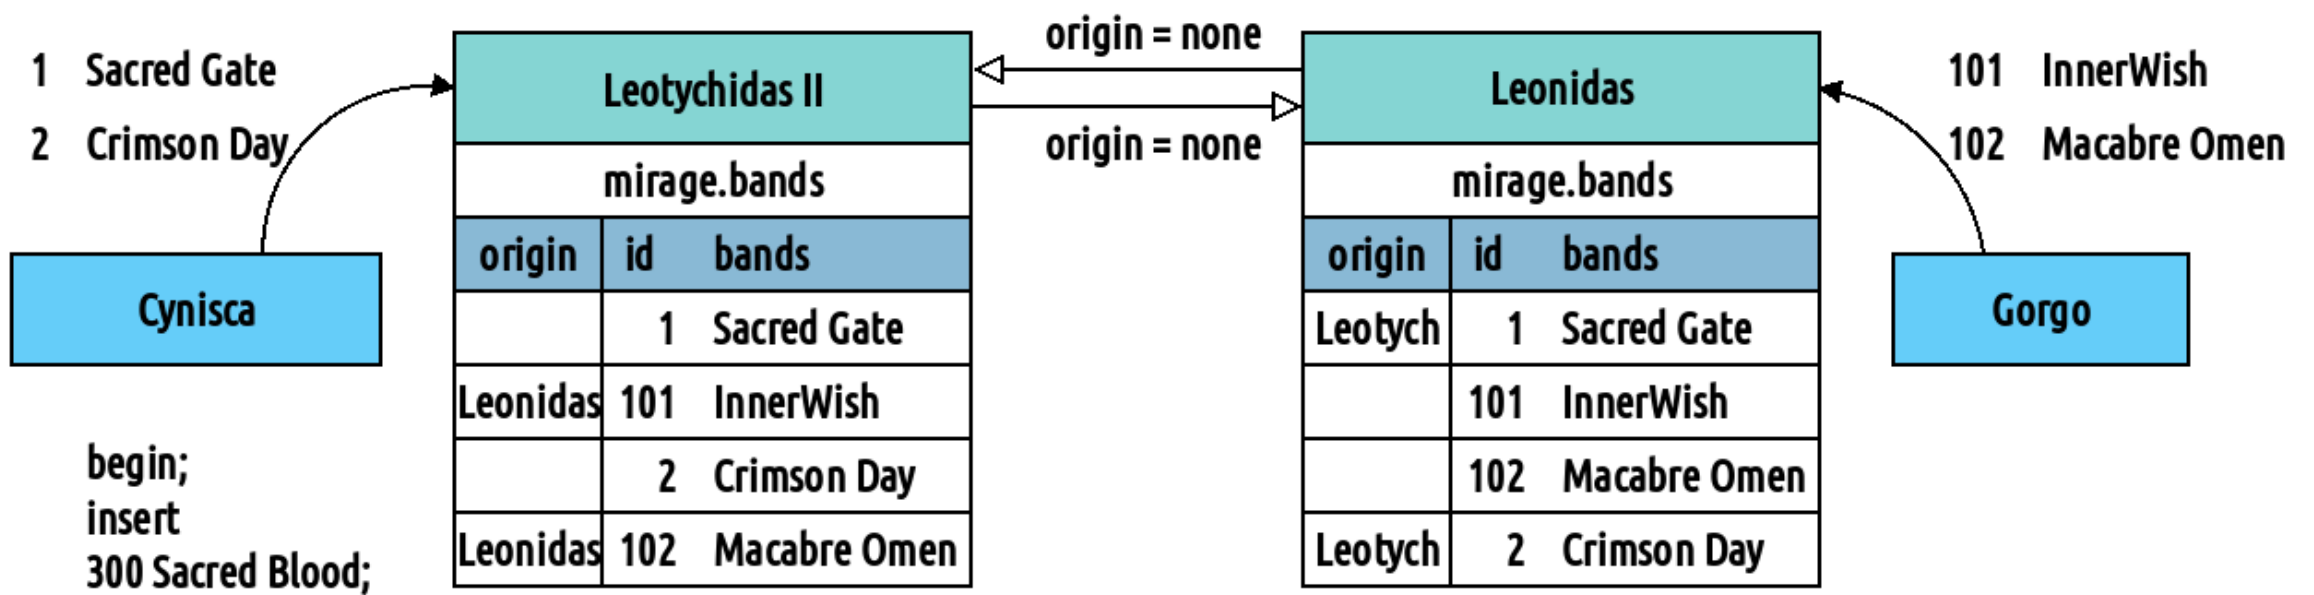
PostgreSQL
cares about the integrity
of your data

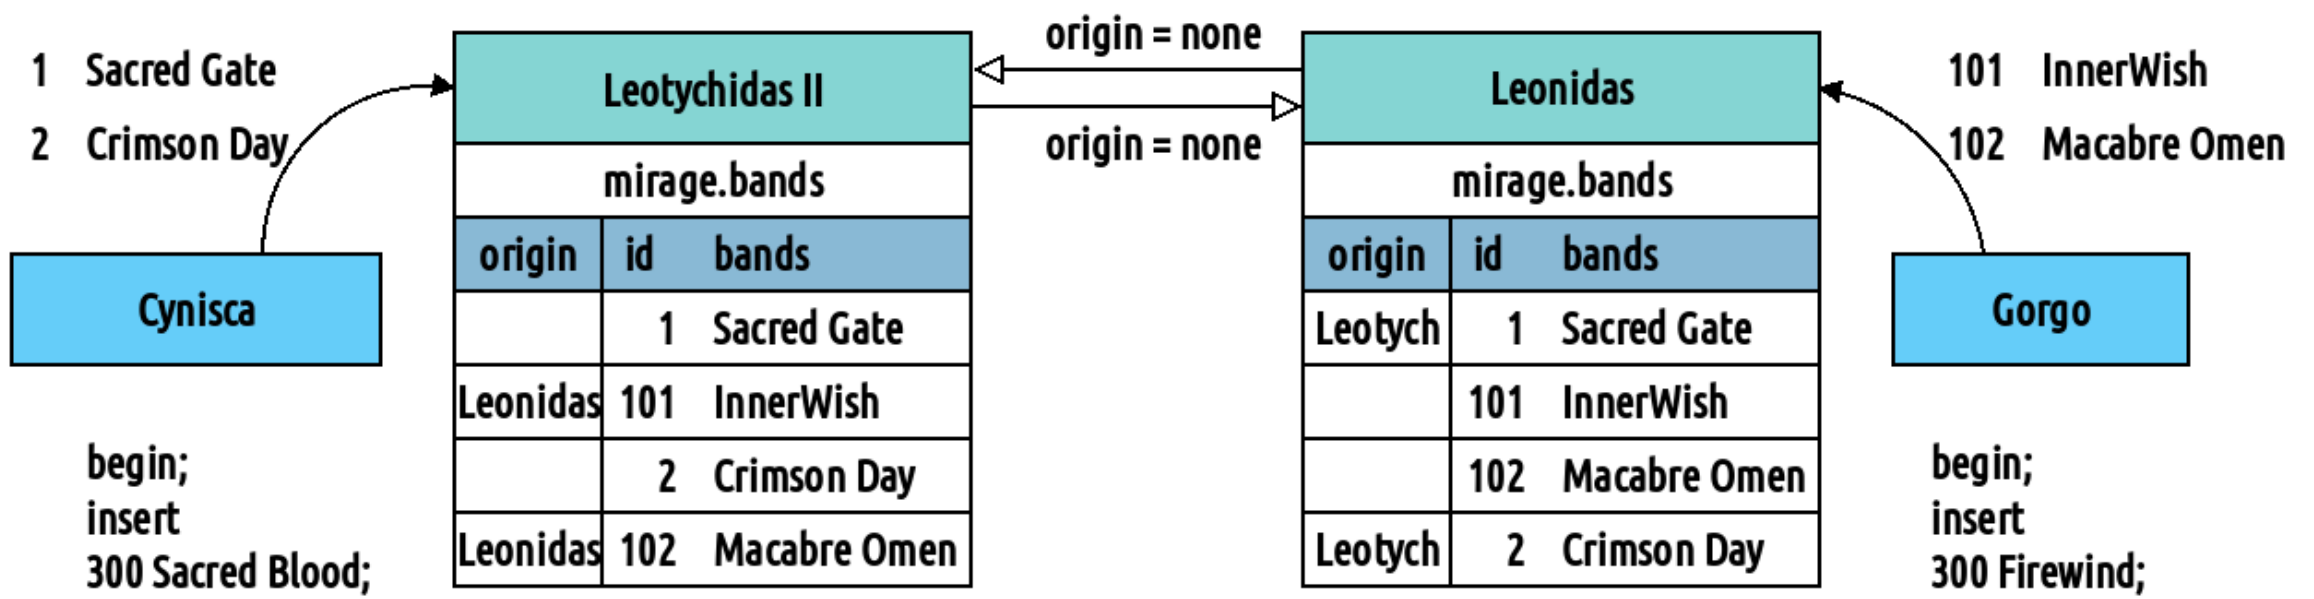


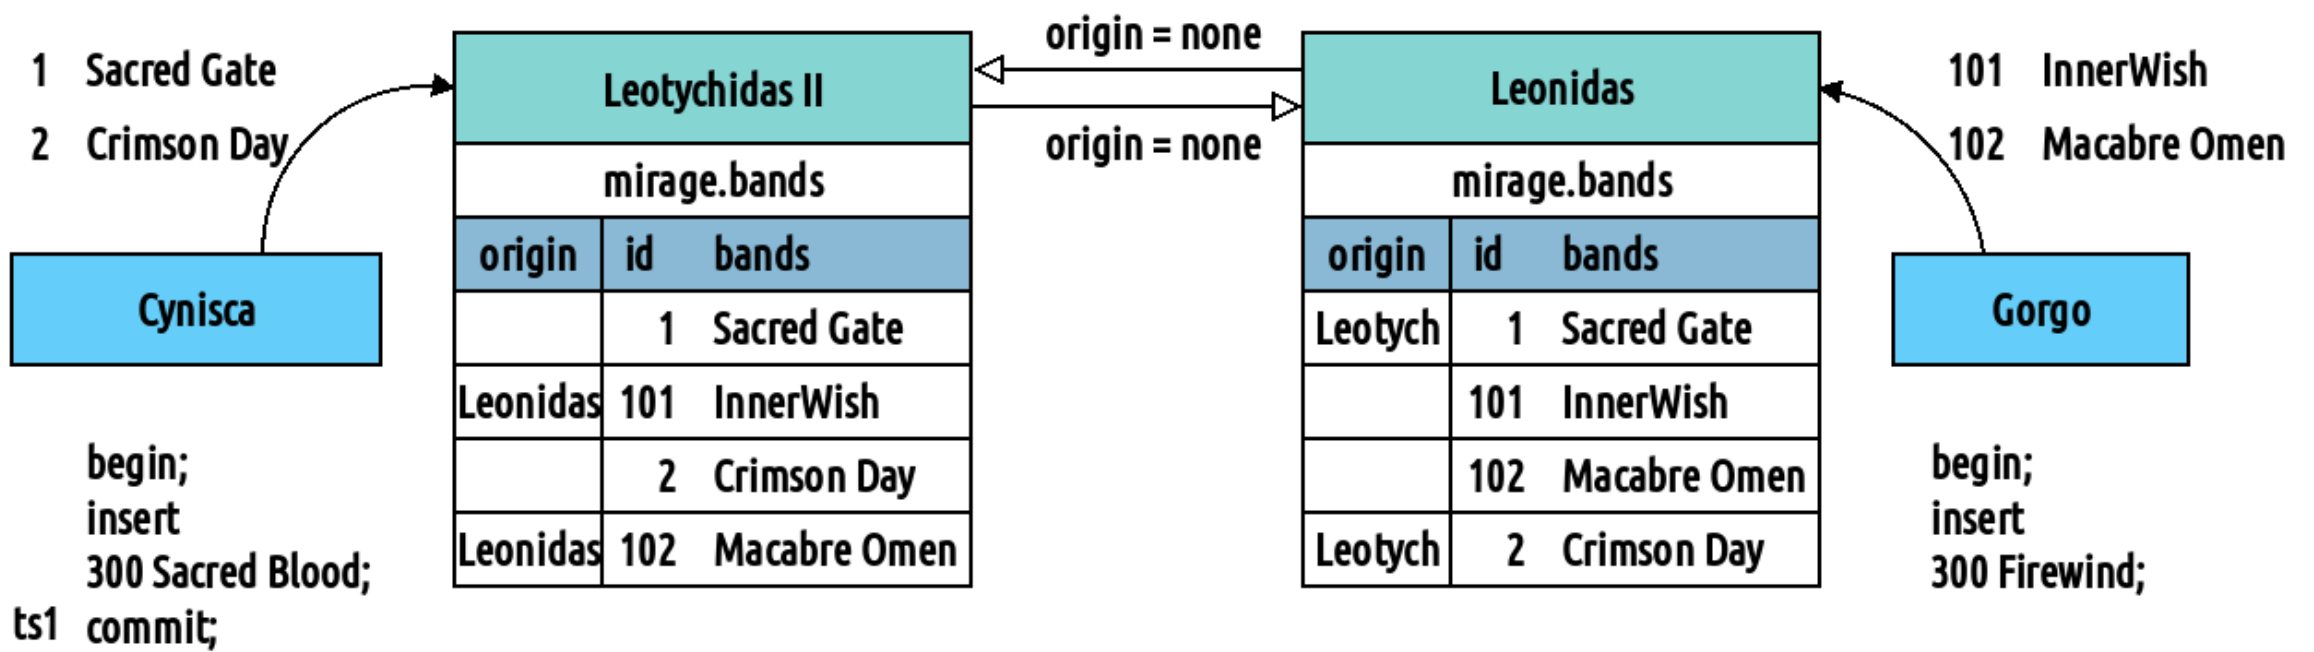
Conflict Resolution

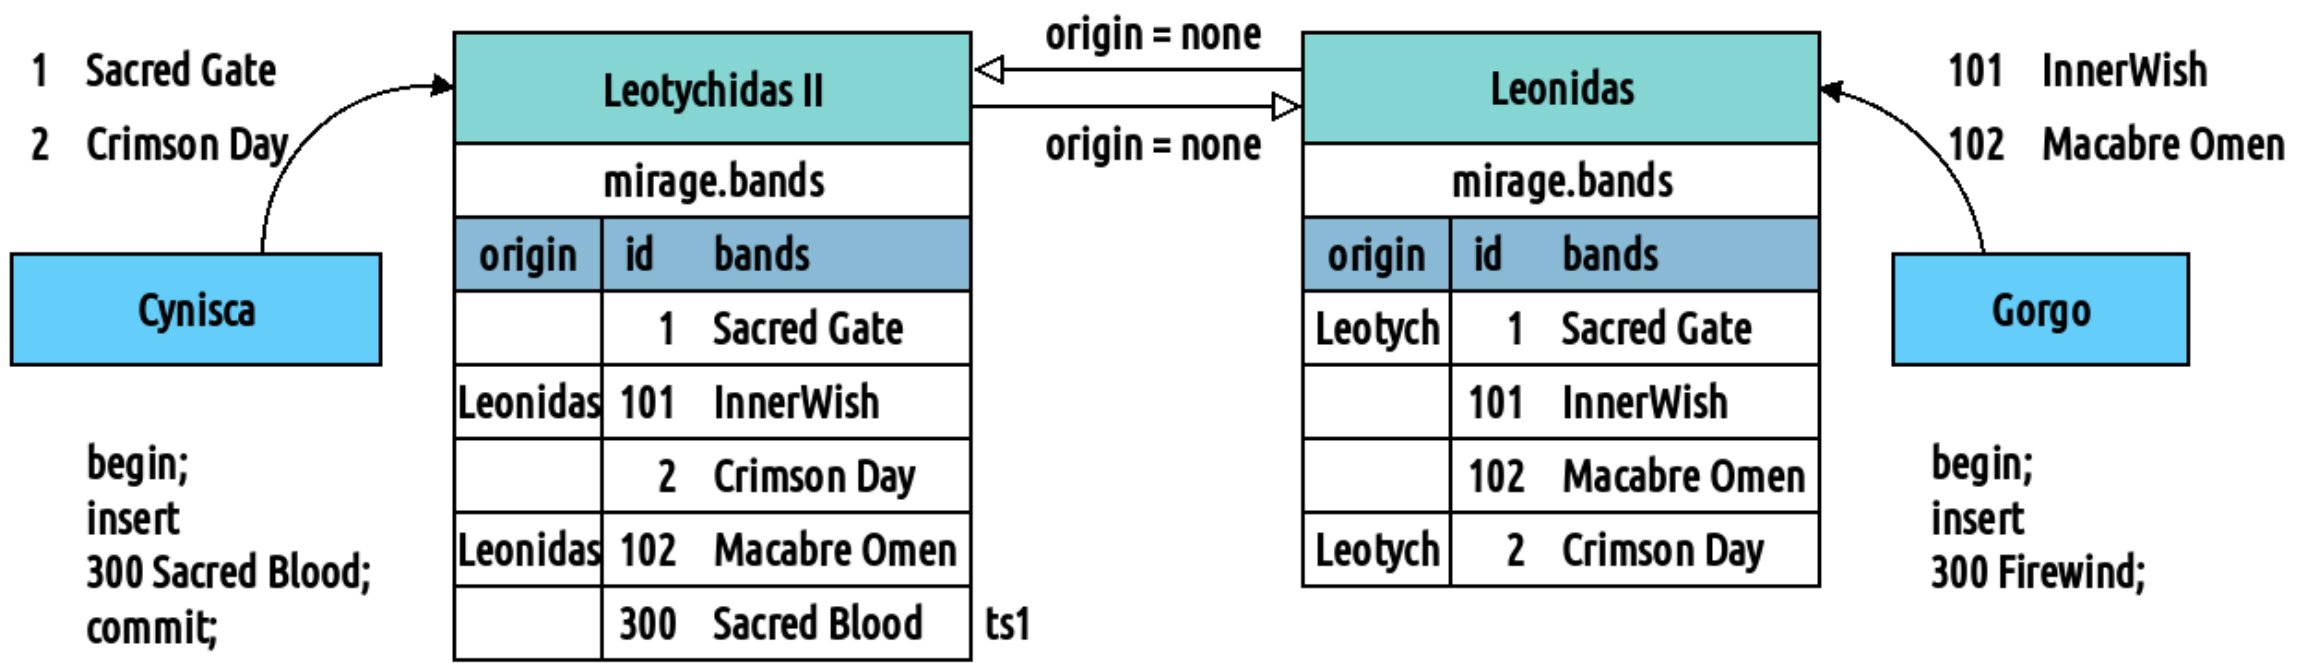


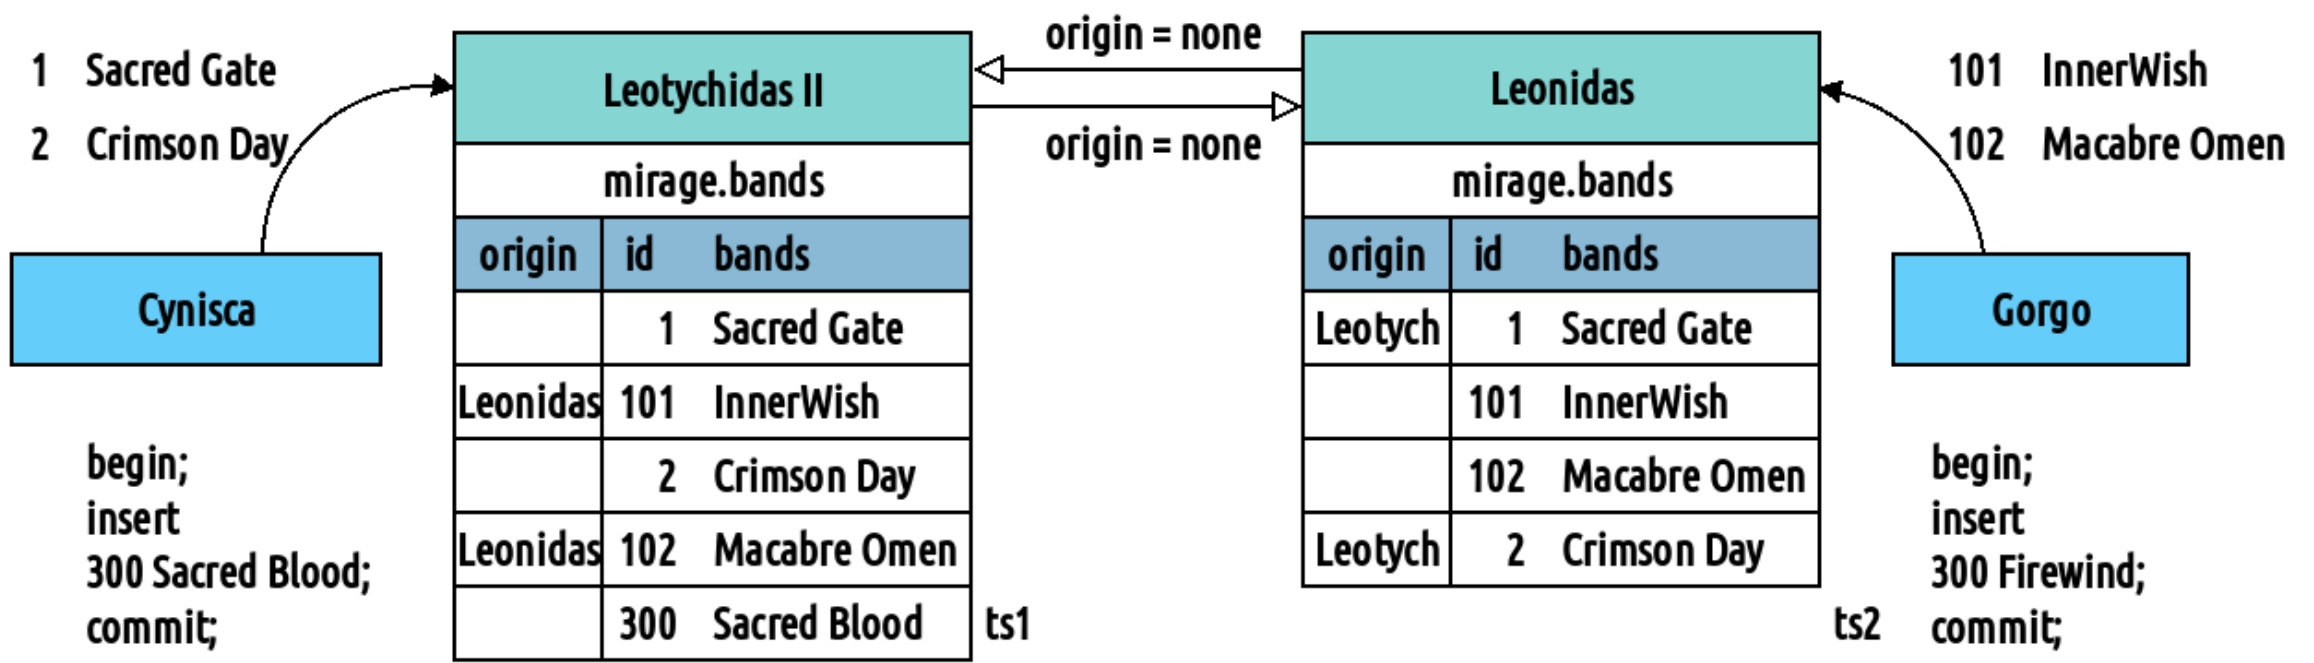





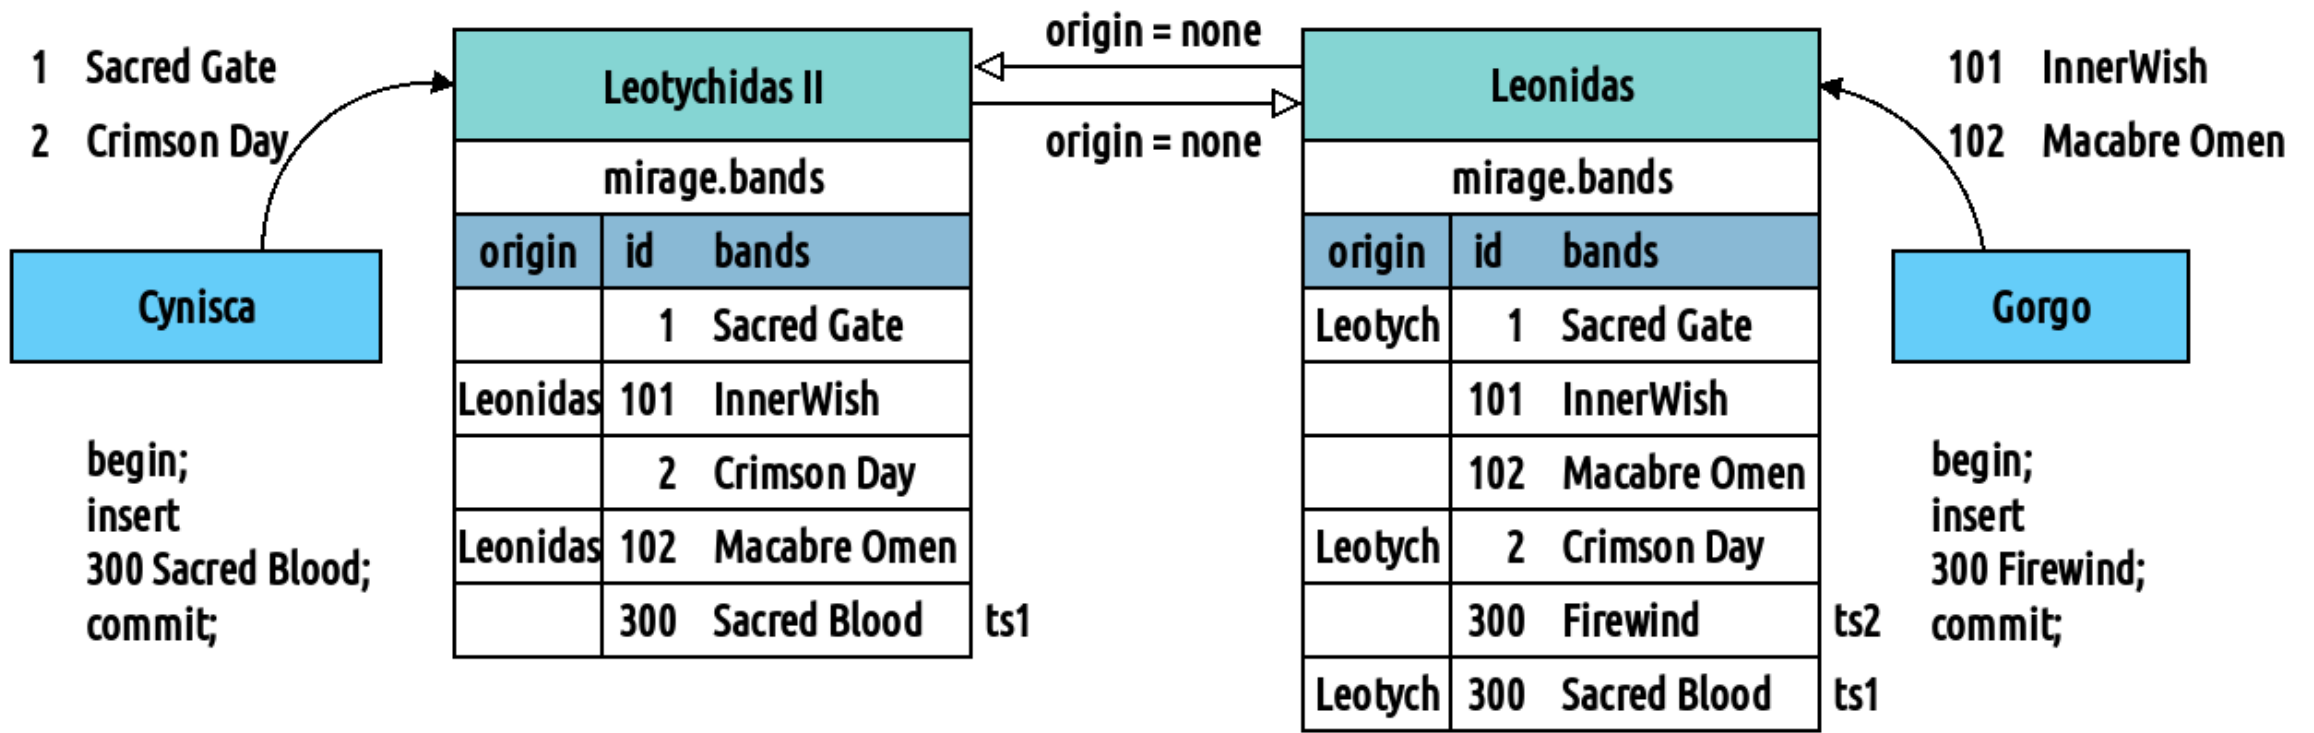


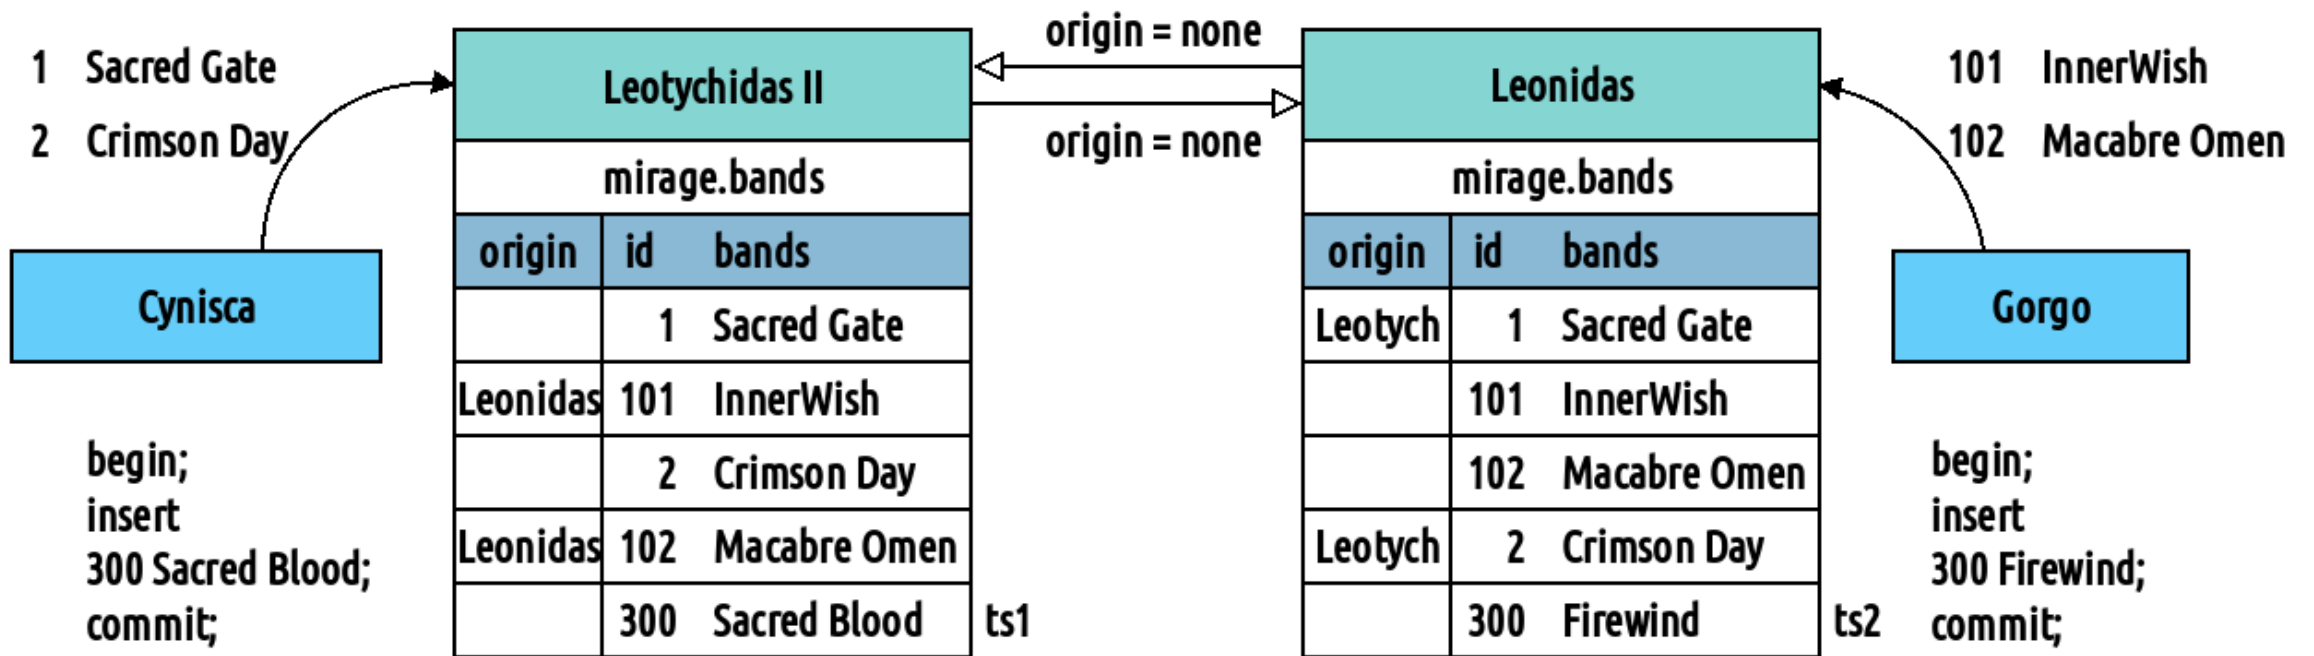




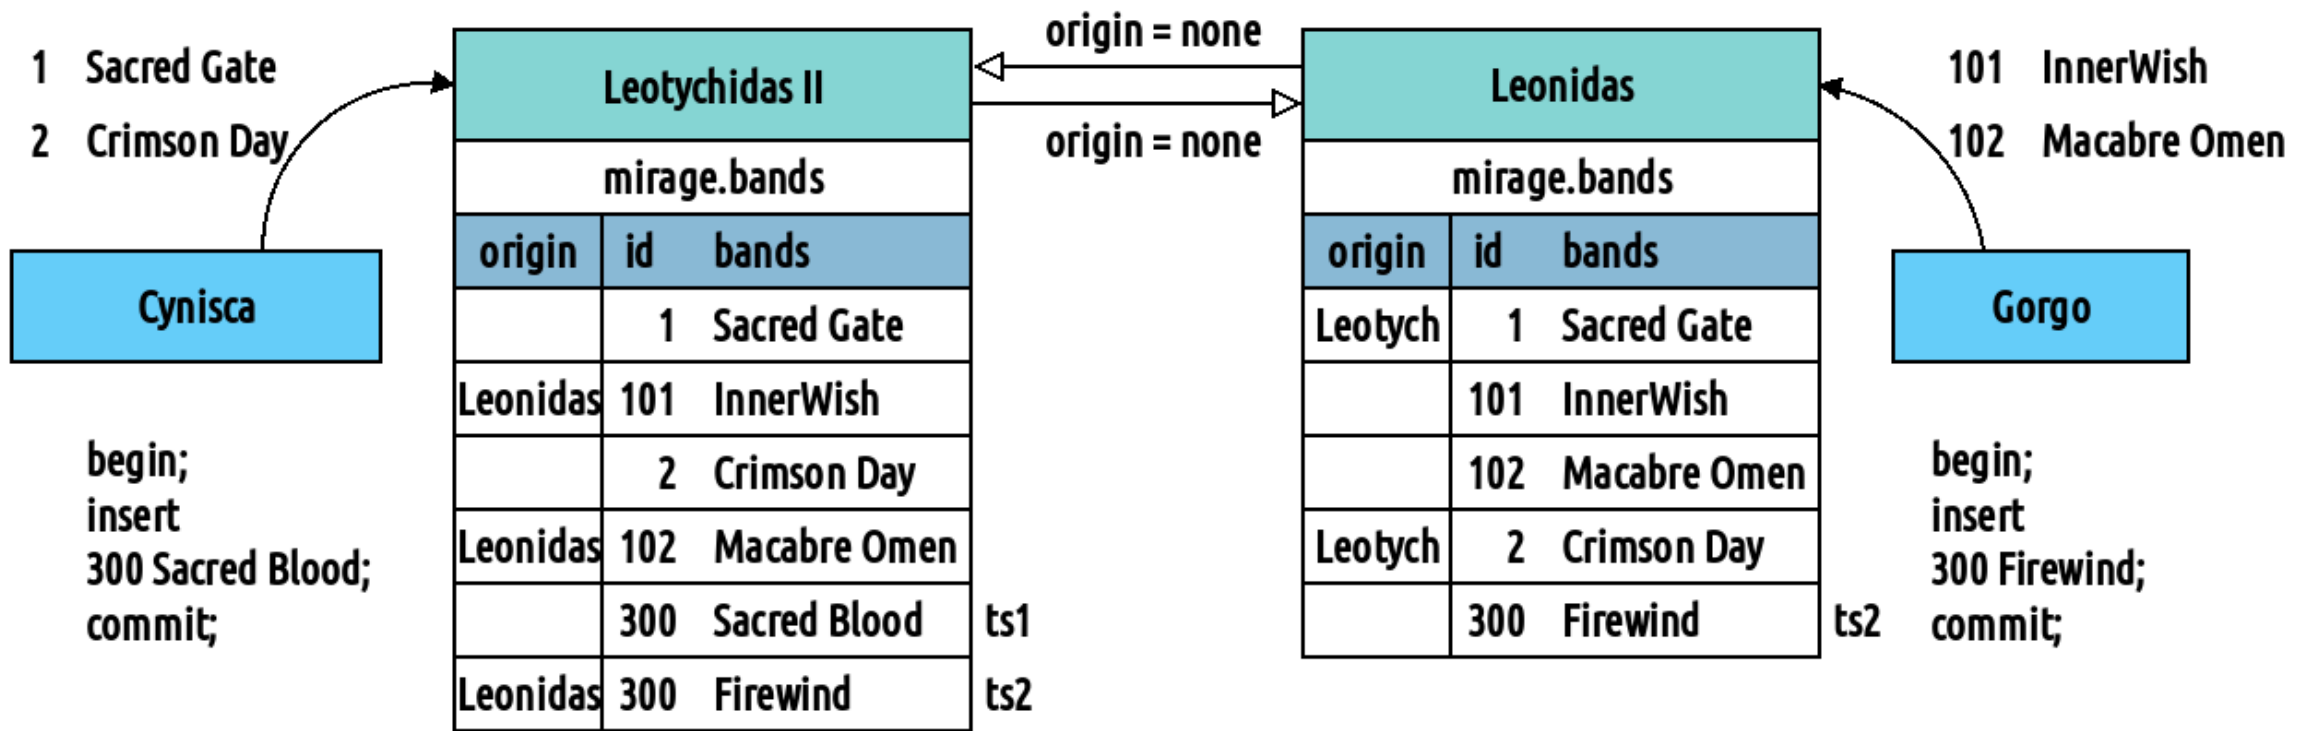




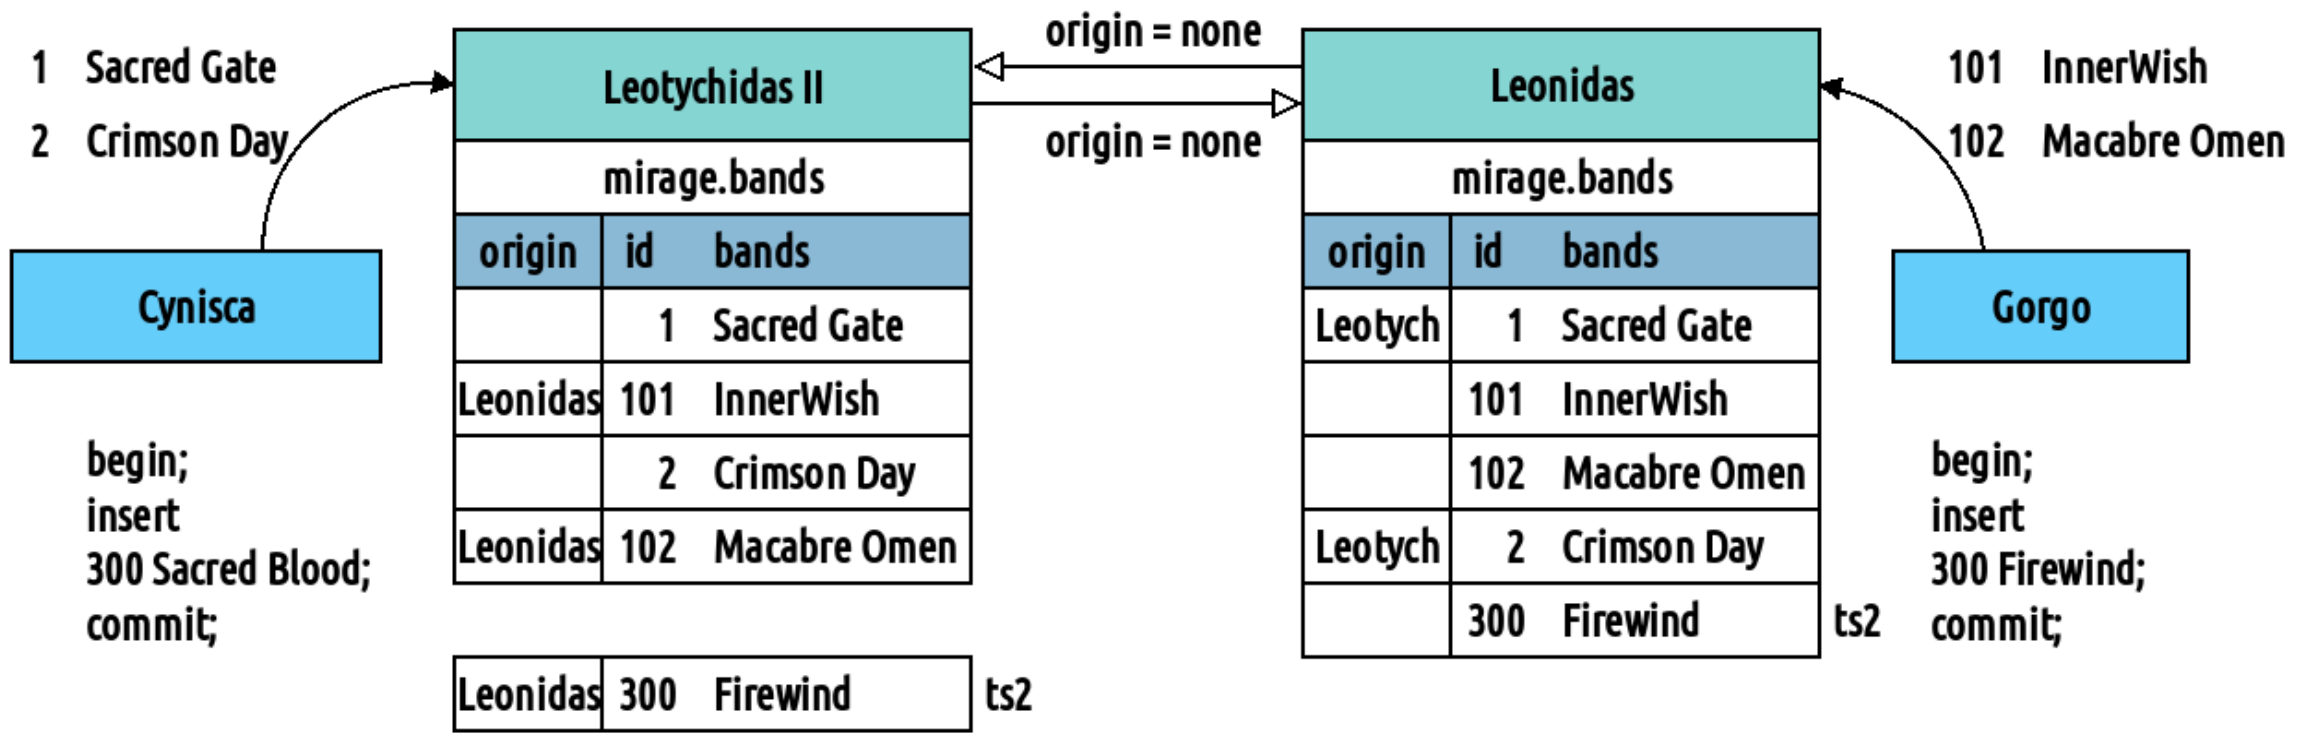




Leotyich	300	Sacred Blood	ts1
----------	-----	--------------	-----

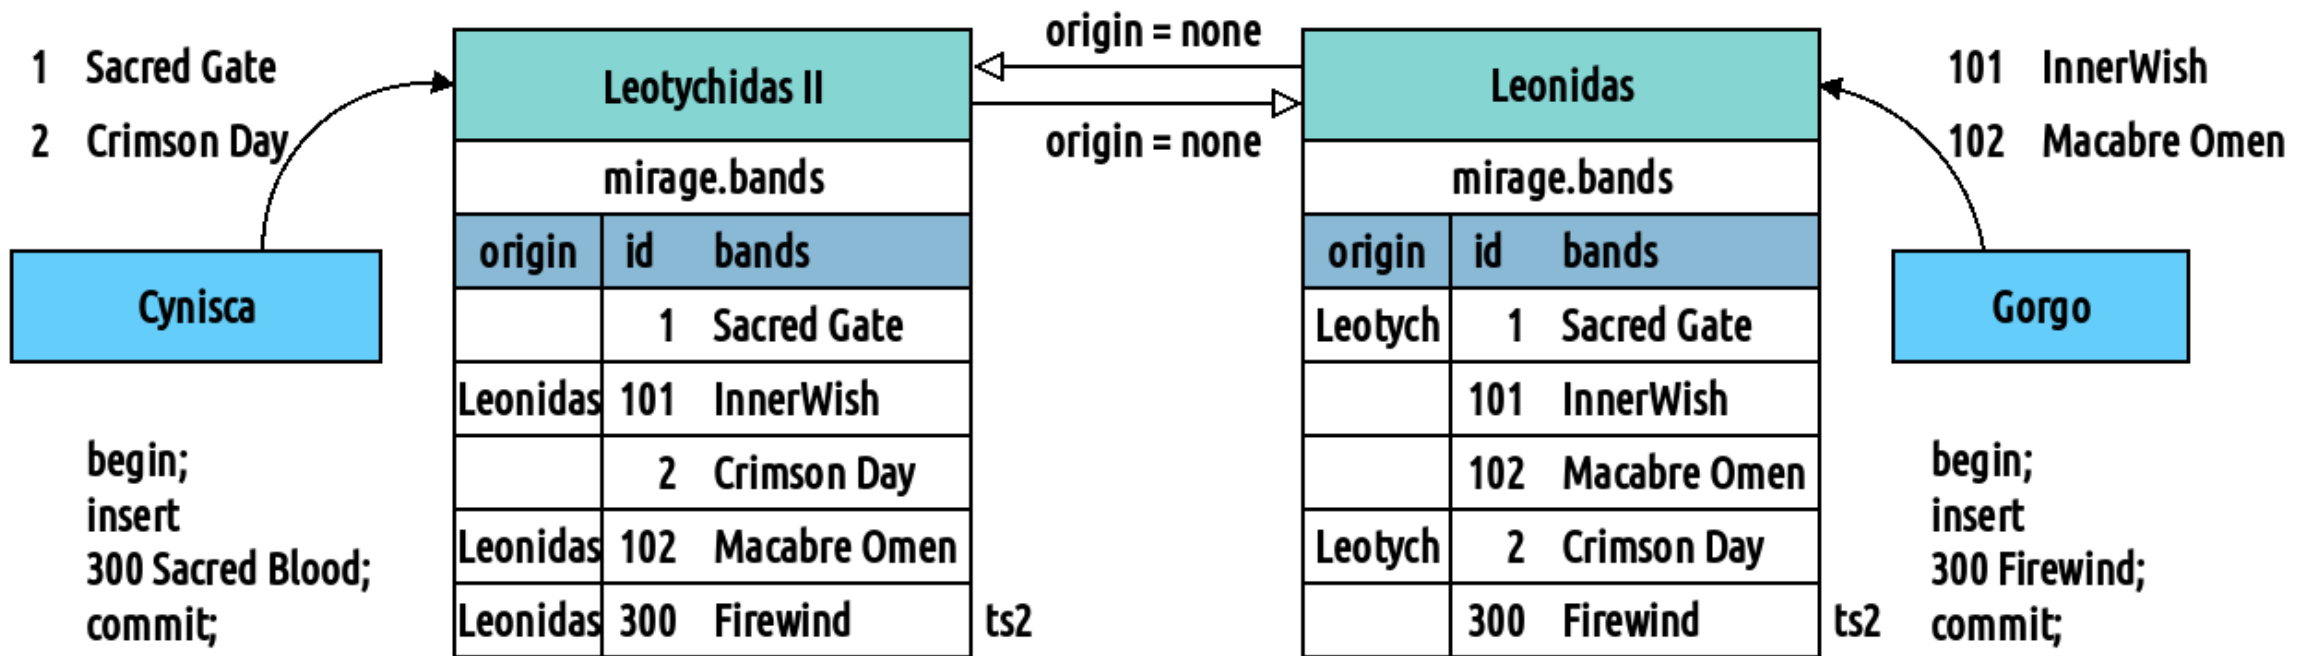


Leotych	300	Sacred Blood	ts1
---------	-----	--------------	-----



	300	Sacred Blood	ts1
--	-----	--------------	-----

Leotych	300	Sacred Blood	ts1
---------	-----	--------------	-----



	300	Sacred Blood	ts1
--	-----	--------------	-----

Leotych	300	Sacred Blood	ts1
---------	-----	--------------	-----

ACID



Atomicity
Consistency
Isolation
Durability


Consistency
Availability
Partition Tolerance



Consistency

Every read receives the most recent write
or an error







	300	Sacred Blood	ts1
--	-----	--------------	-----

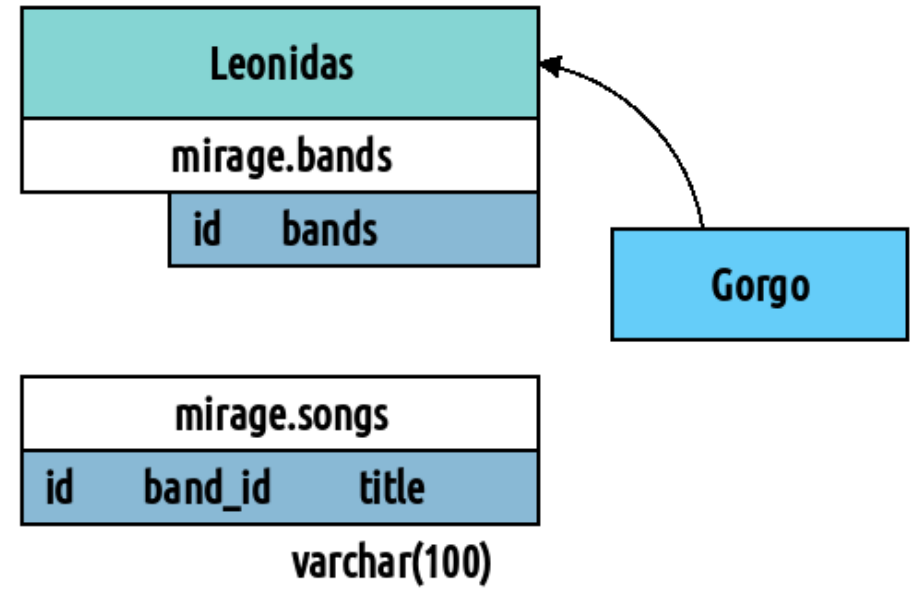
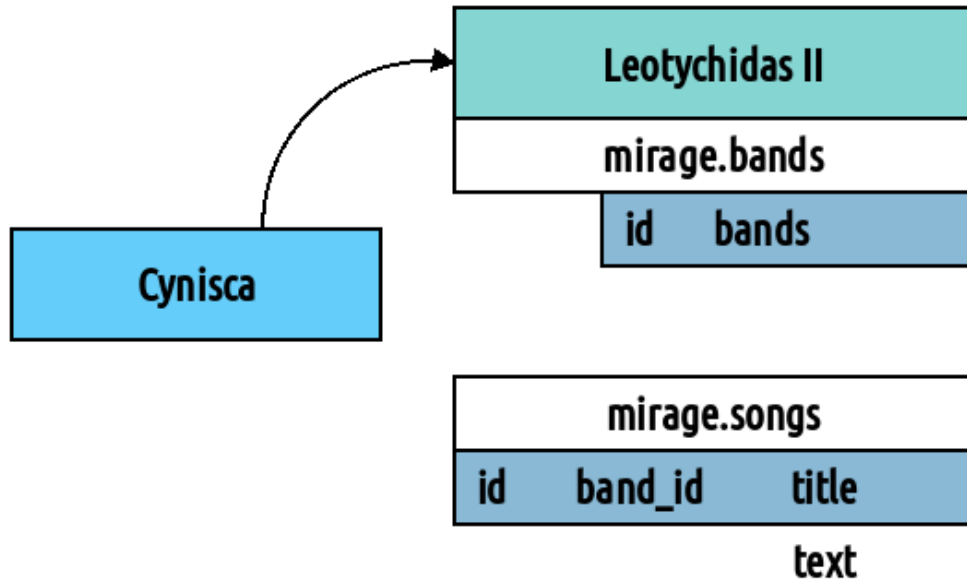
Leotyich	300	Sacred Blood	ts1
----------	-----	--------------	-----

DDL

Data Definition Language







DDL Based on Consensus



Dual Kingship and the Gerousia



Dual Kingship and the Gerousia



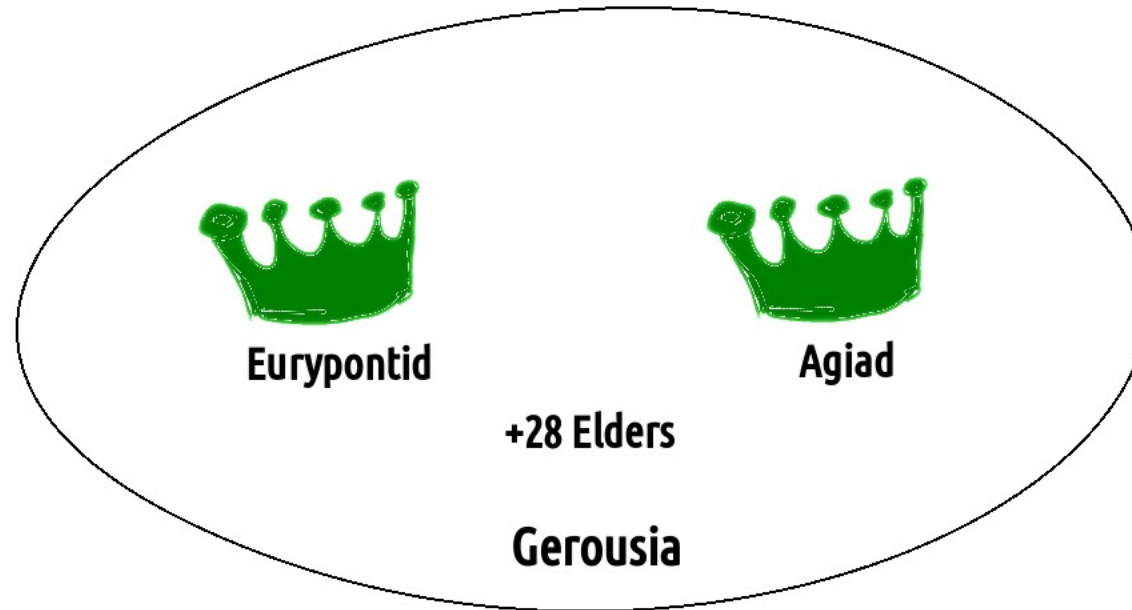
Eurypontid





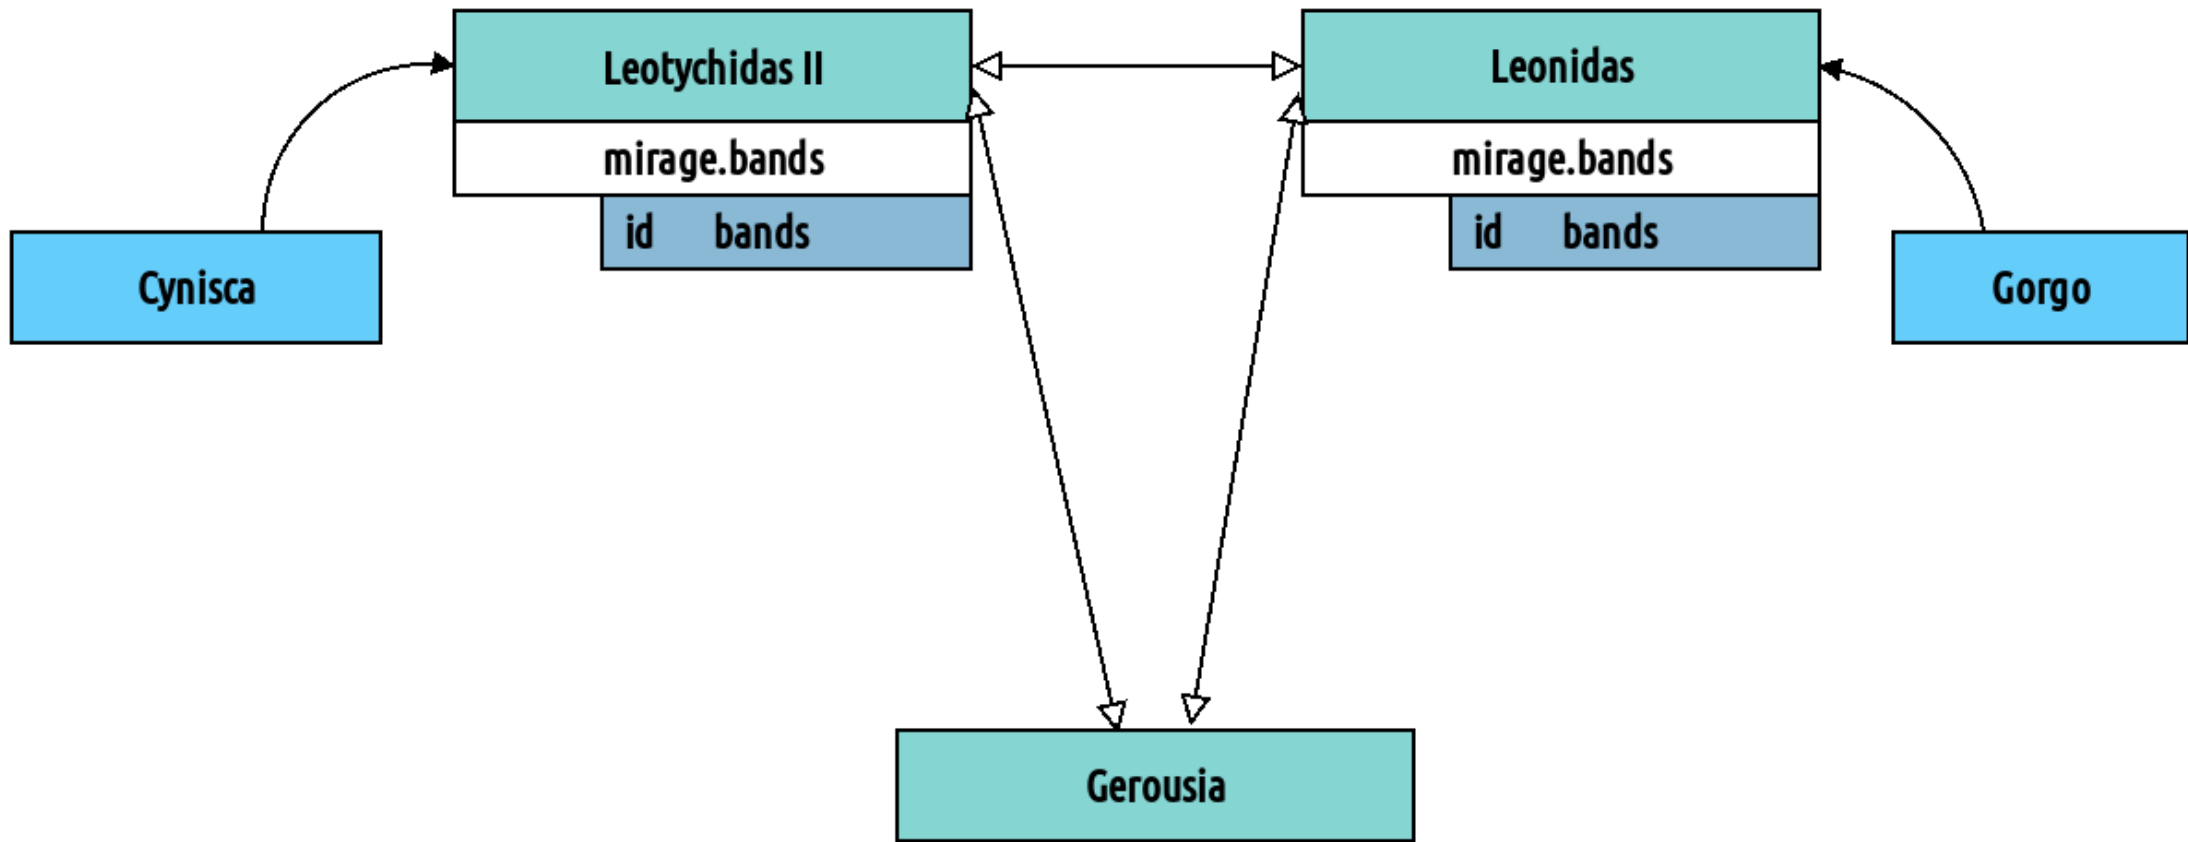
Agiad

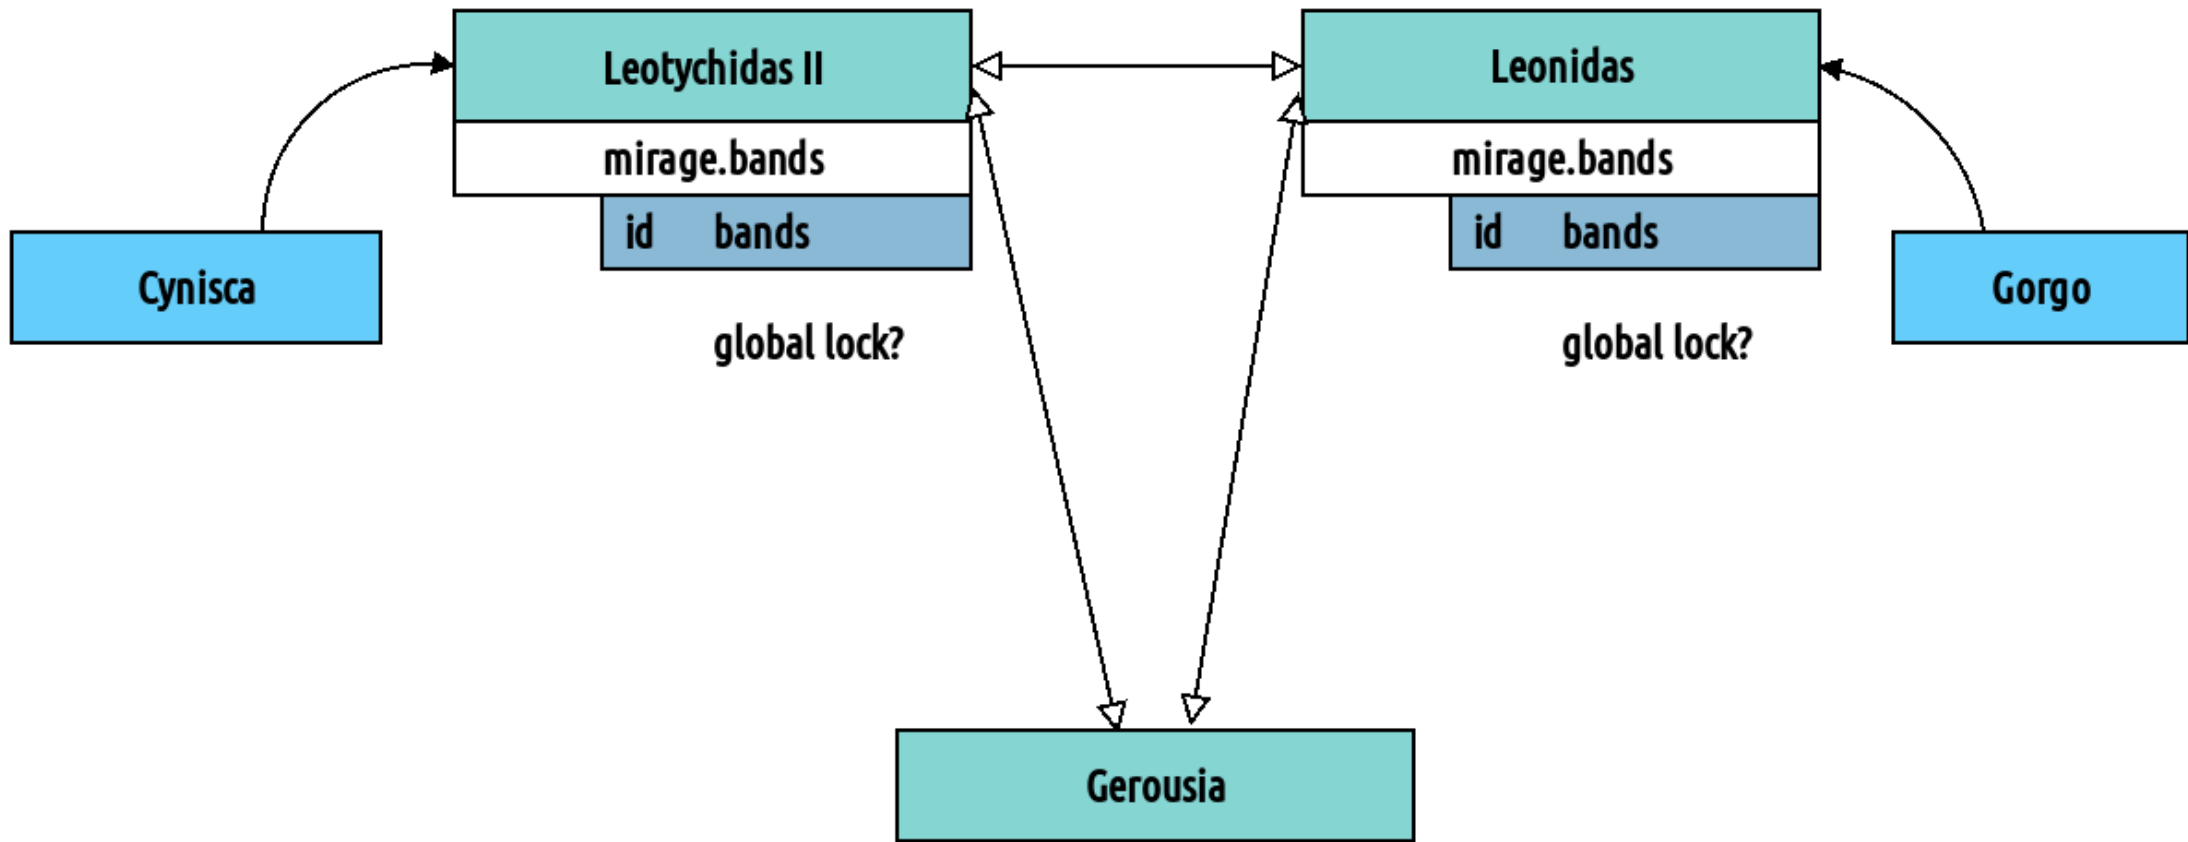


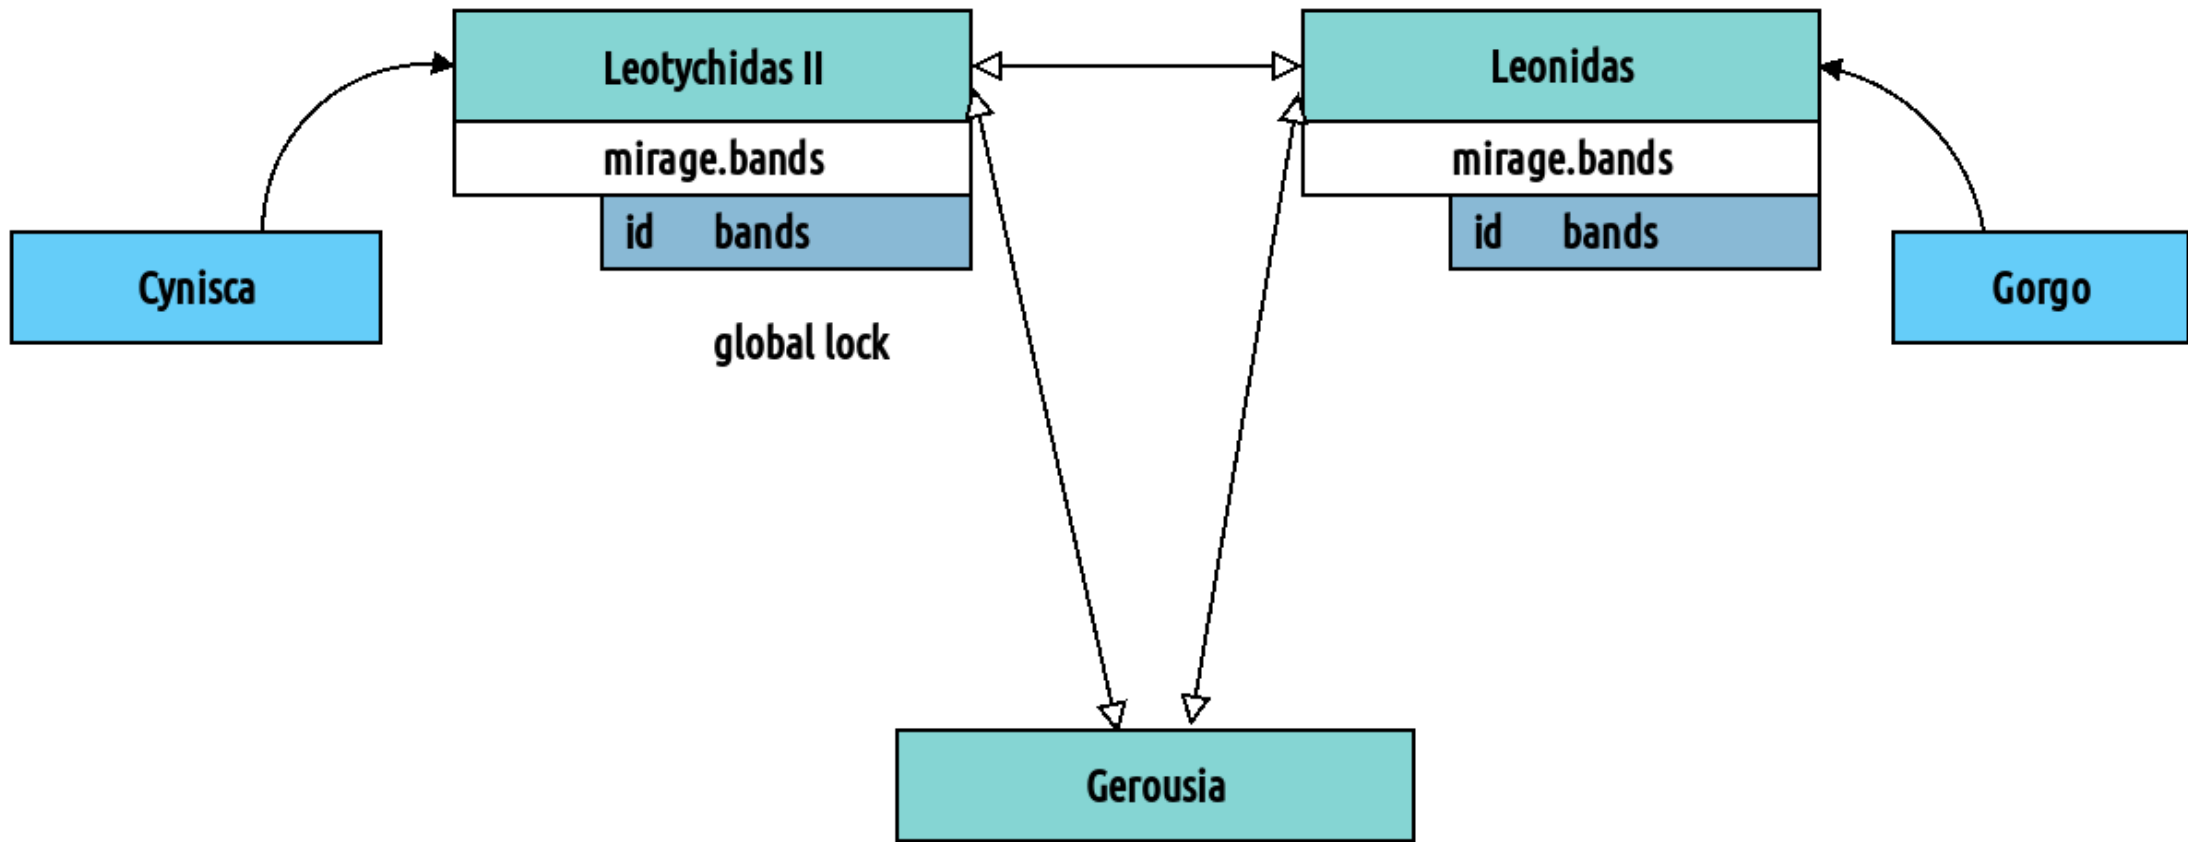
Dual Kingship and the Gerousia

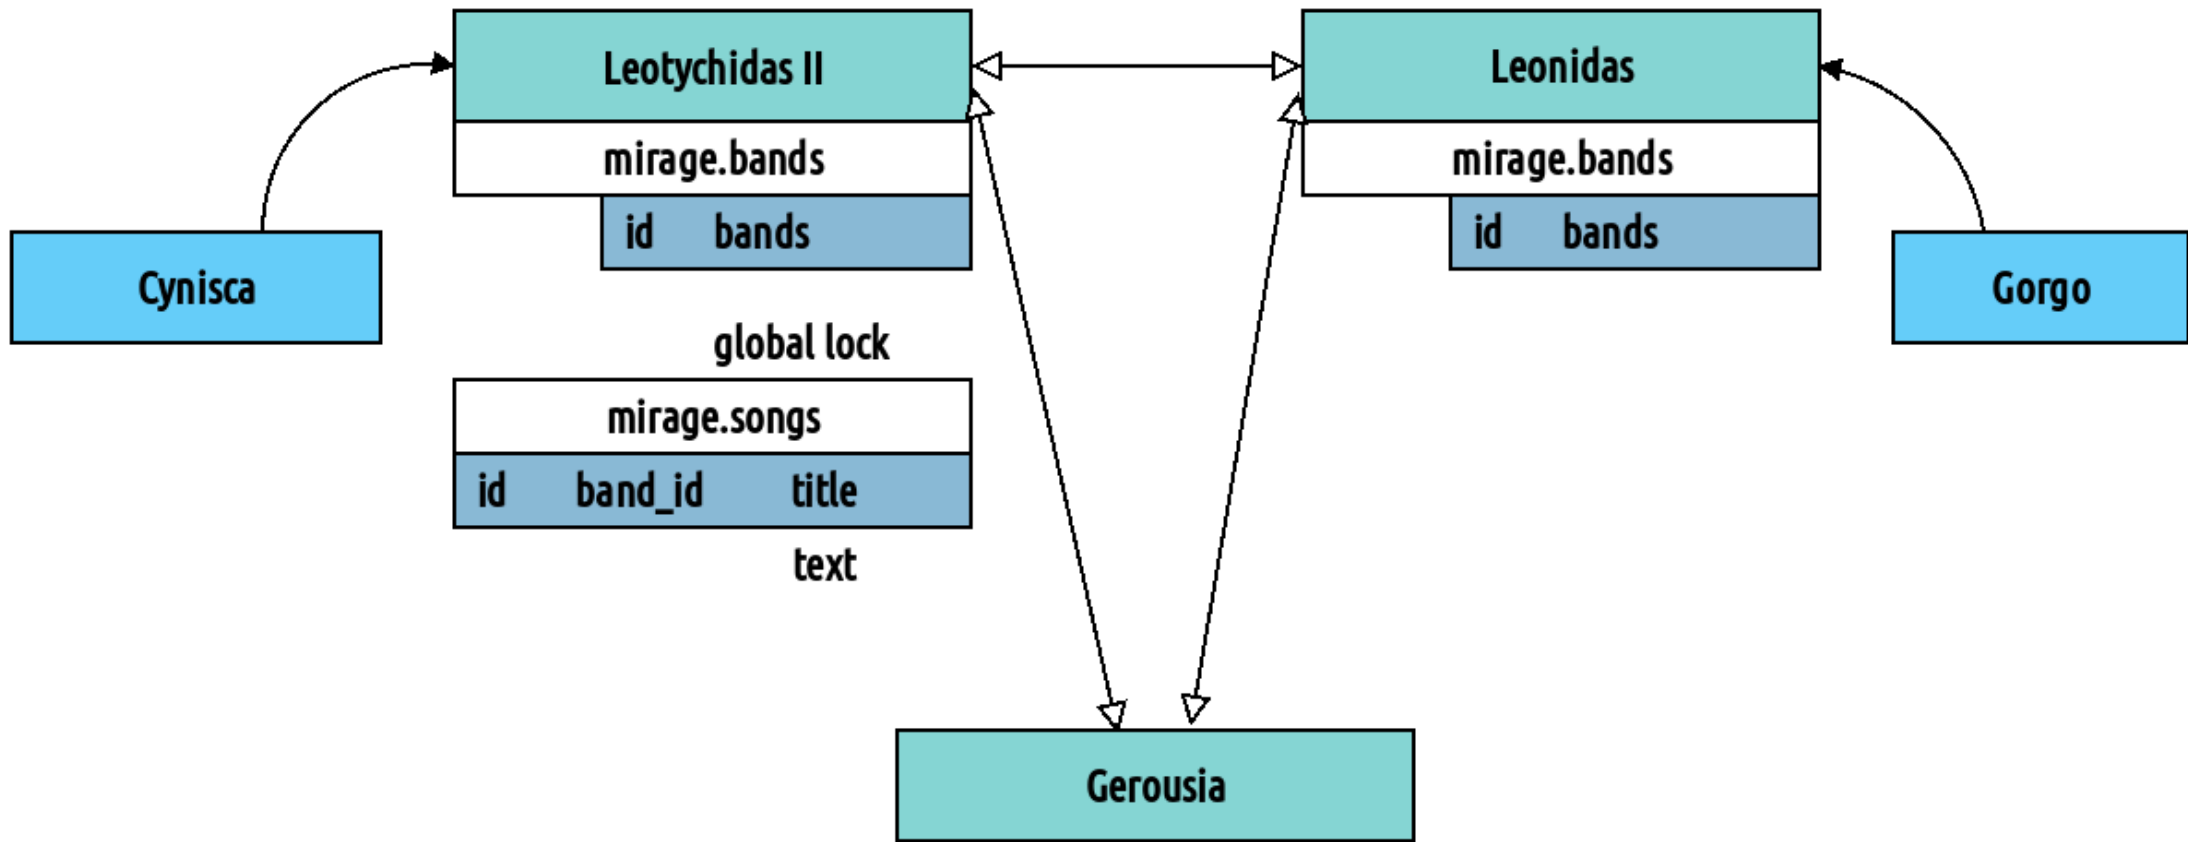


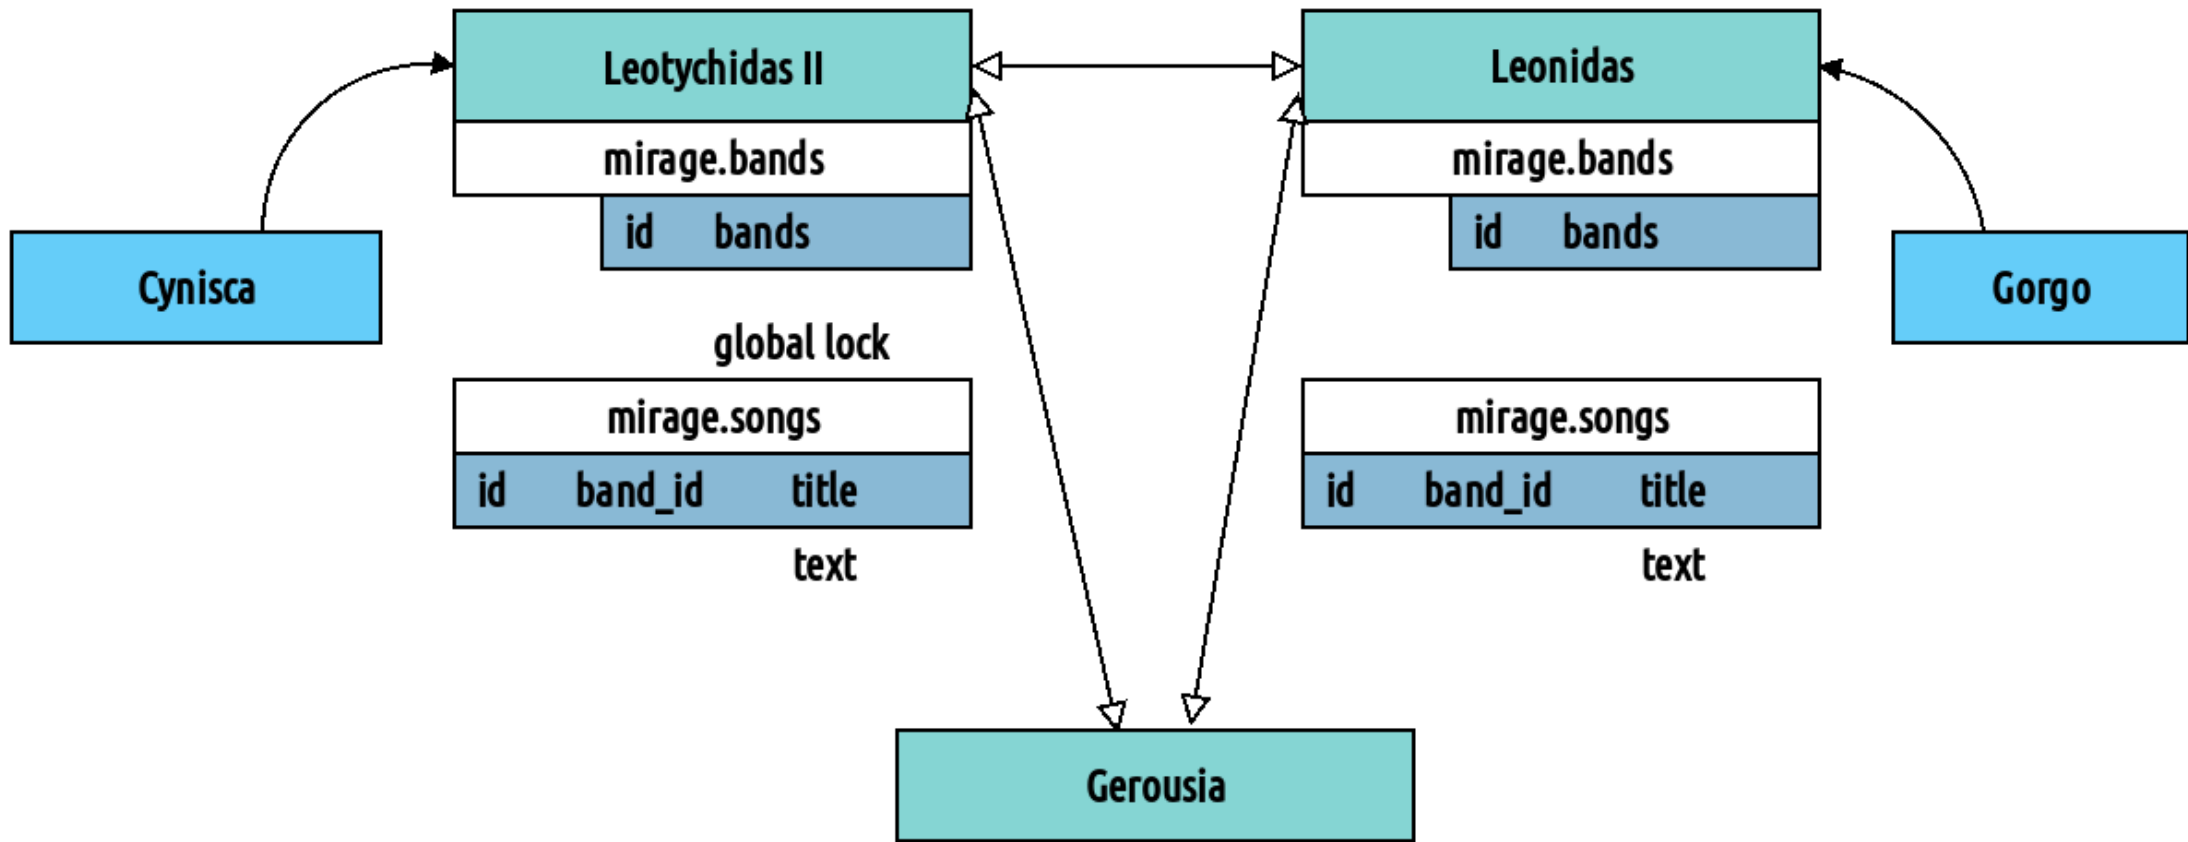















Can we use the Gerousia for something else?






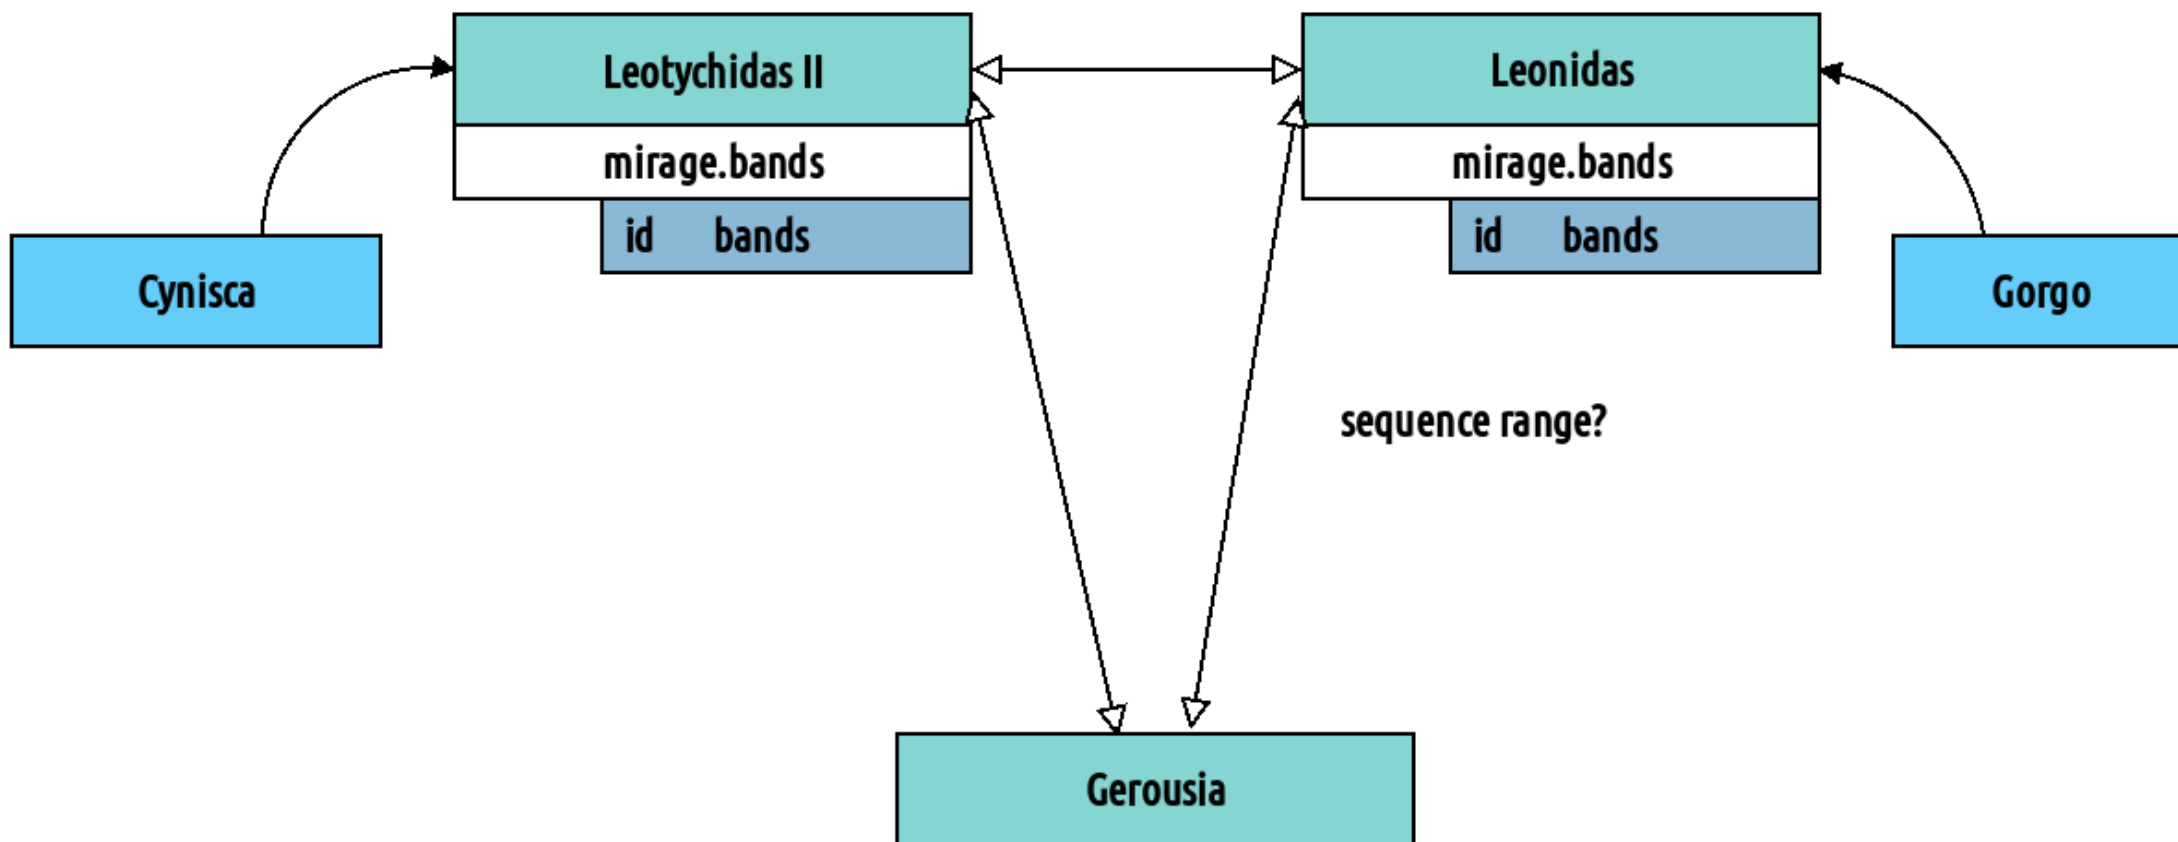
CCC

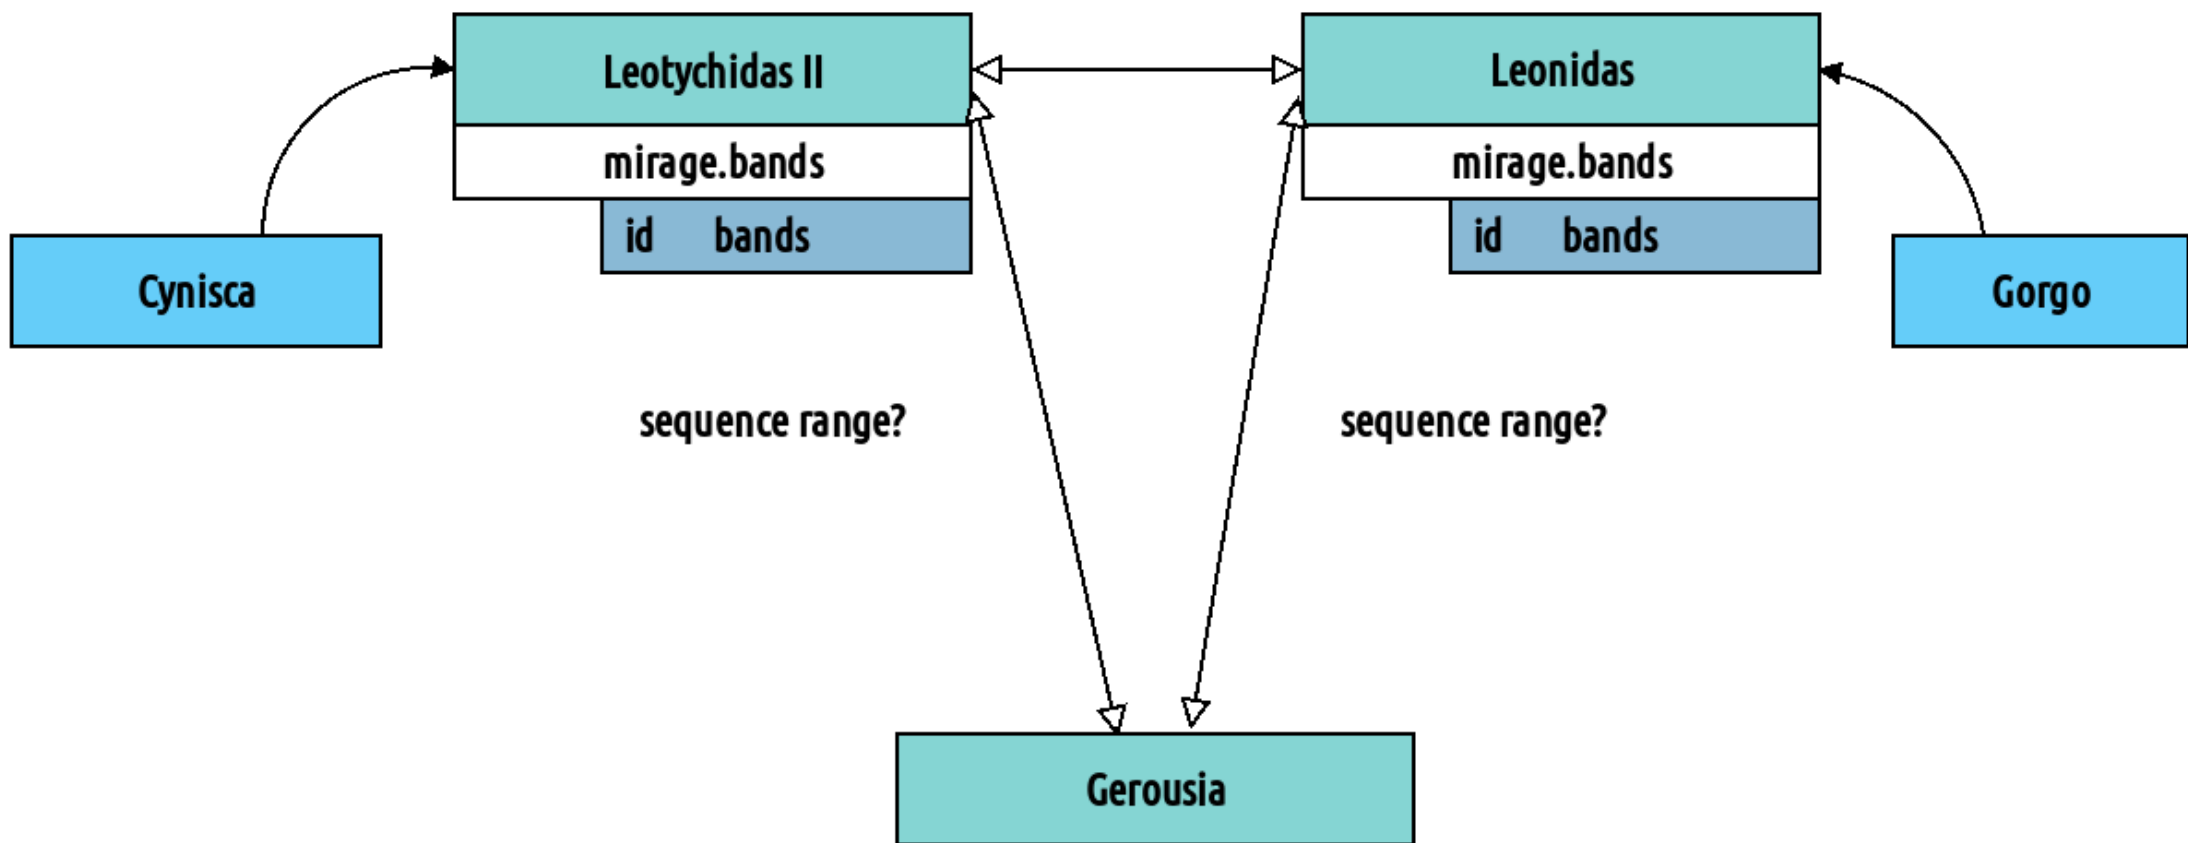


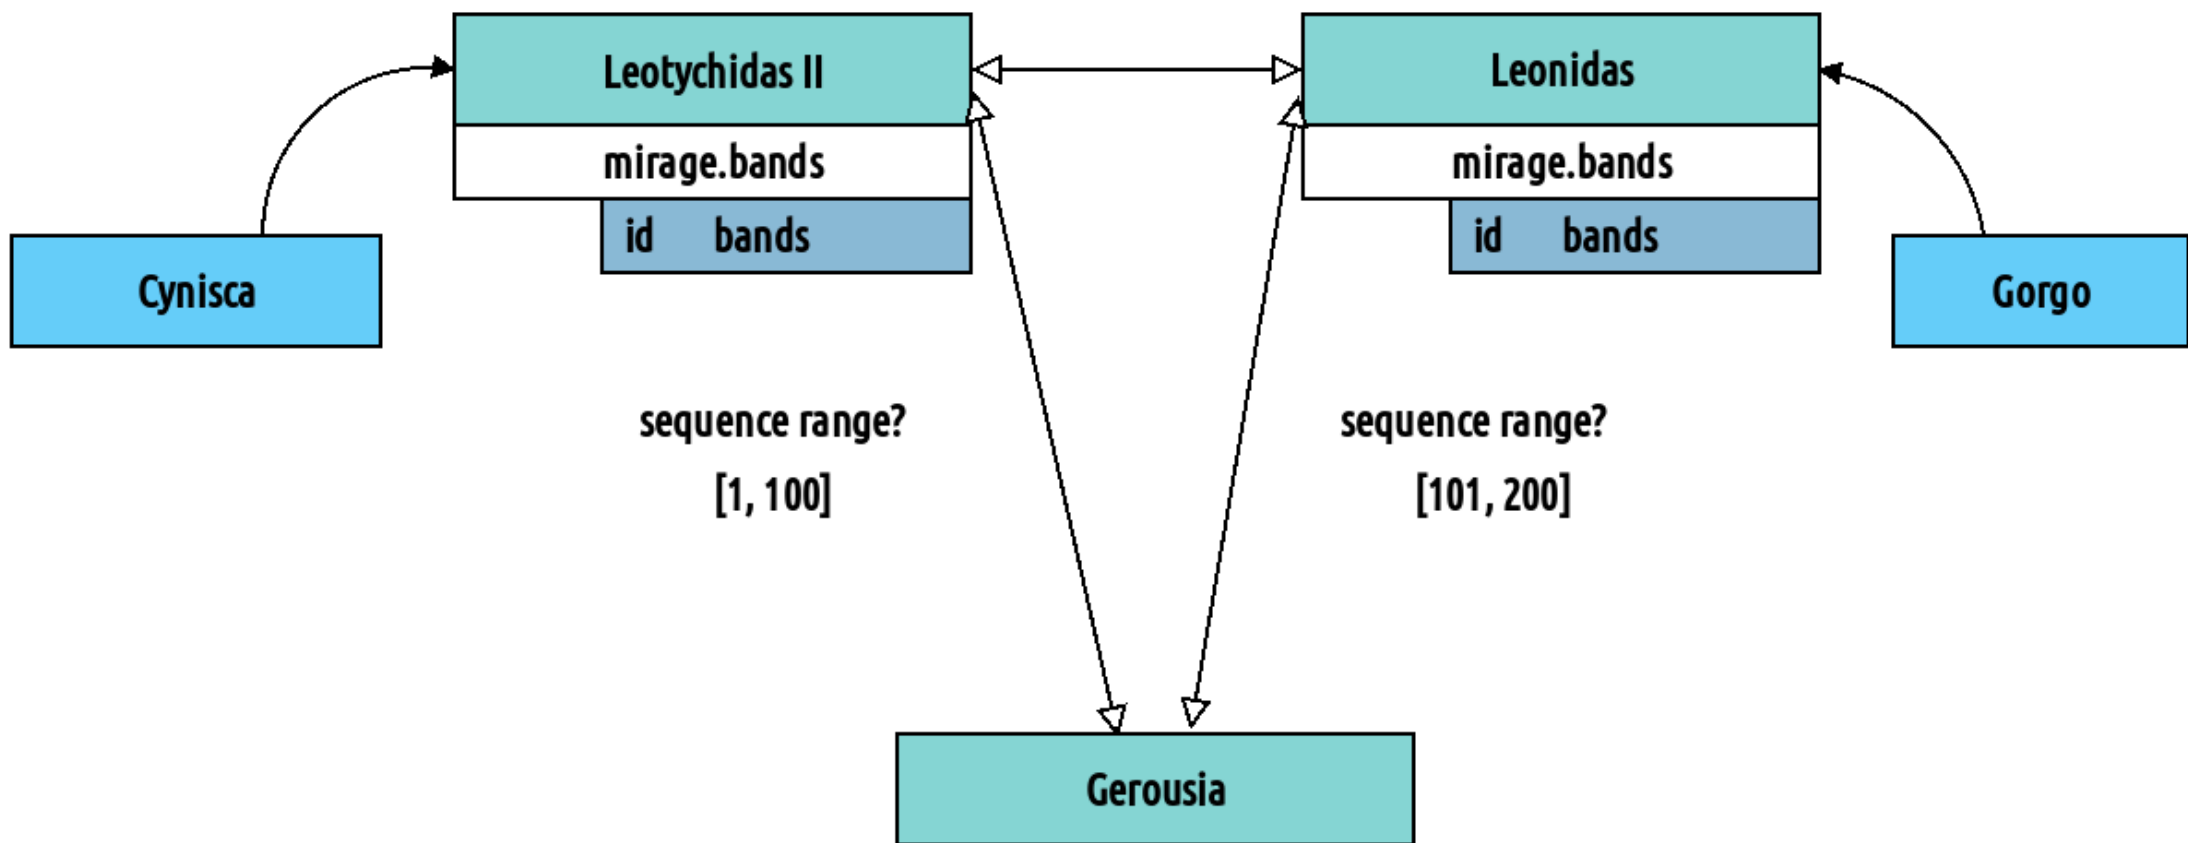
300




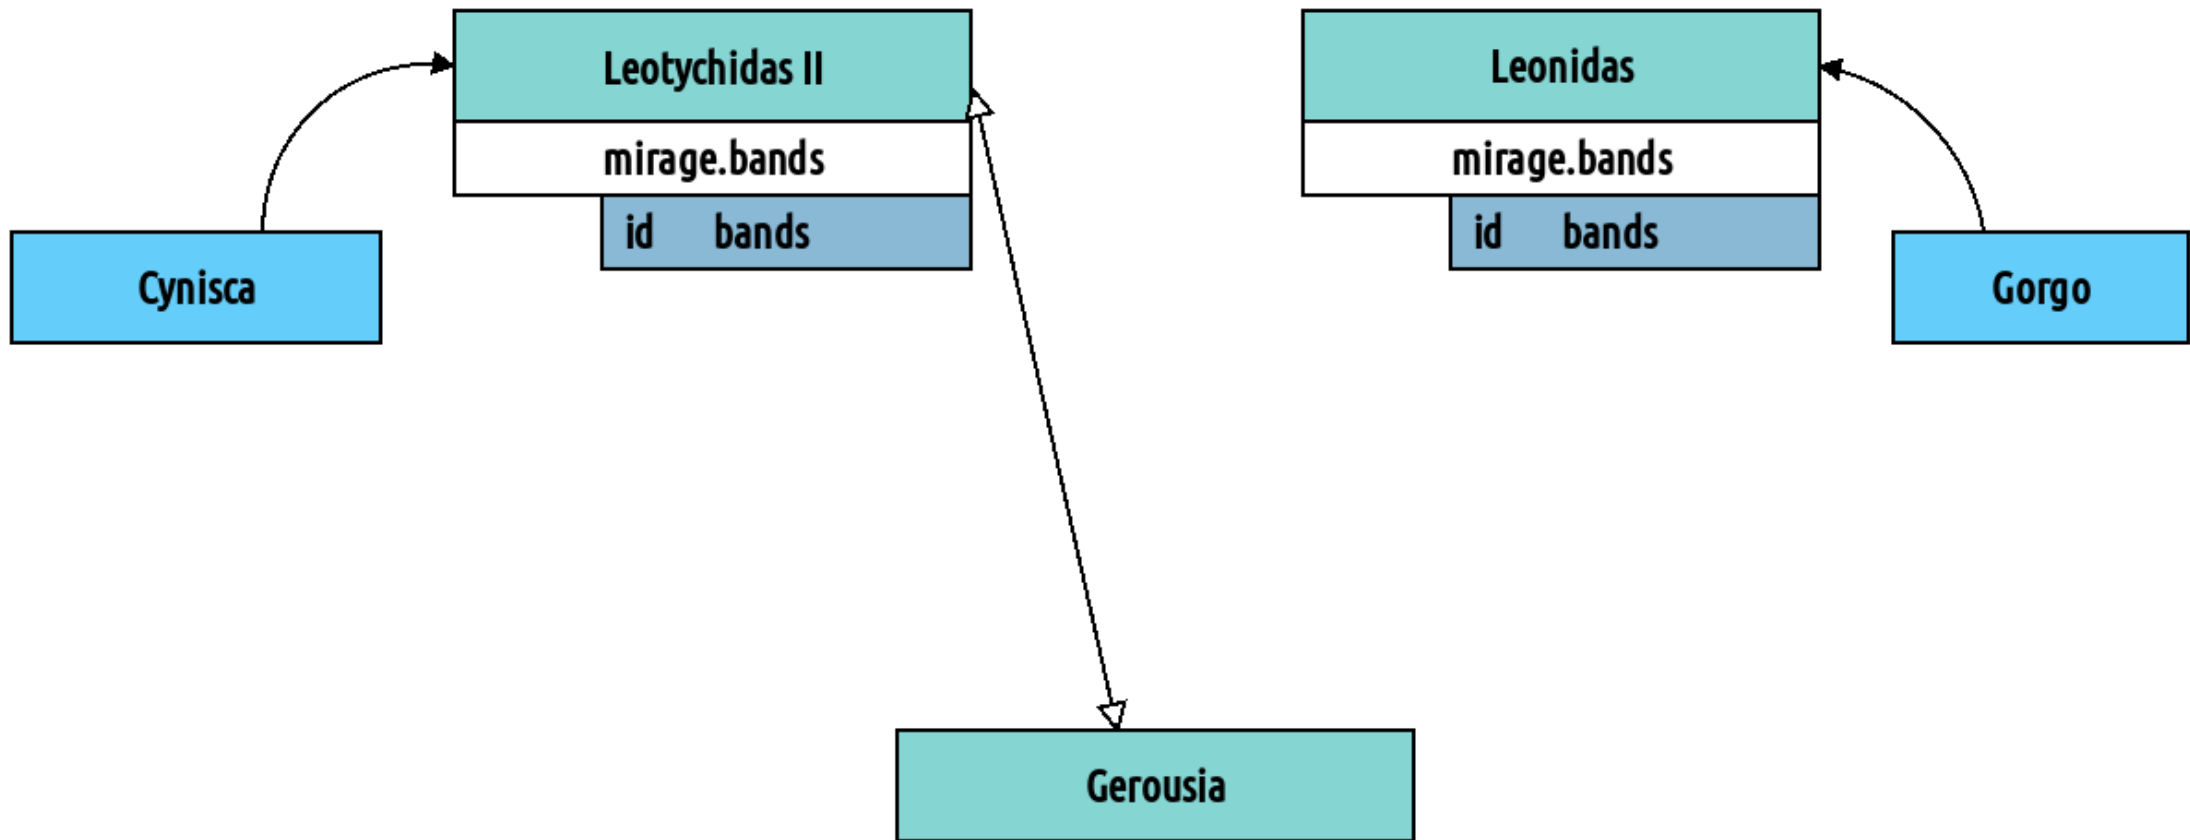


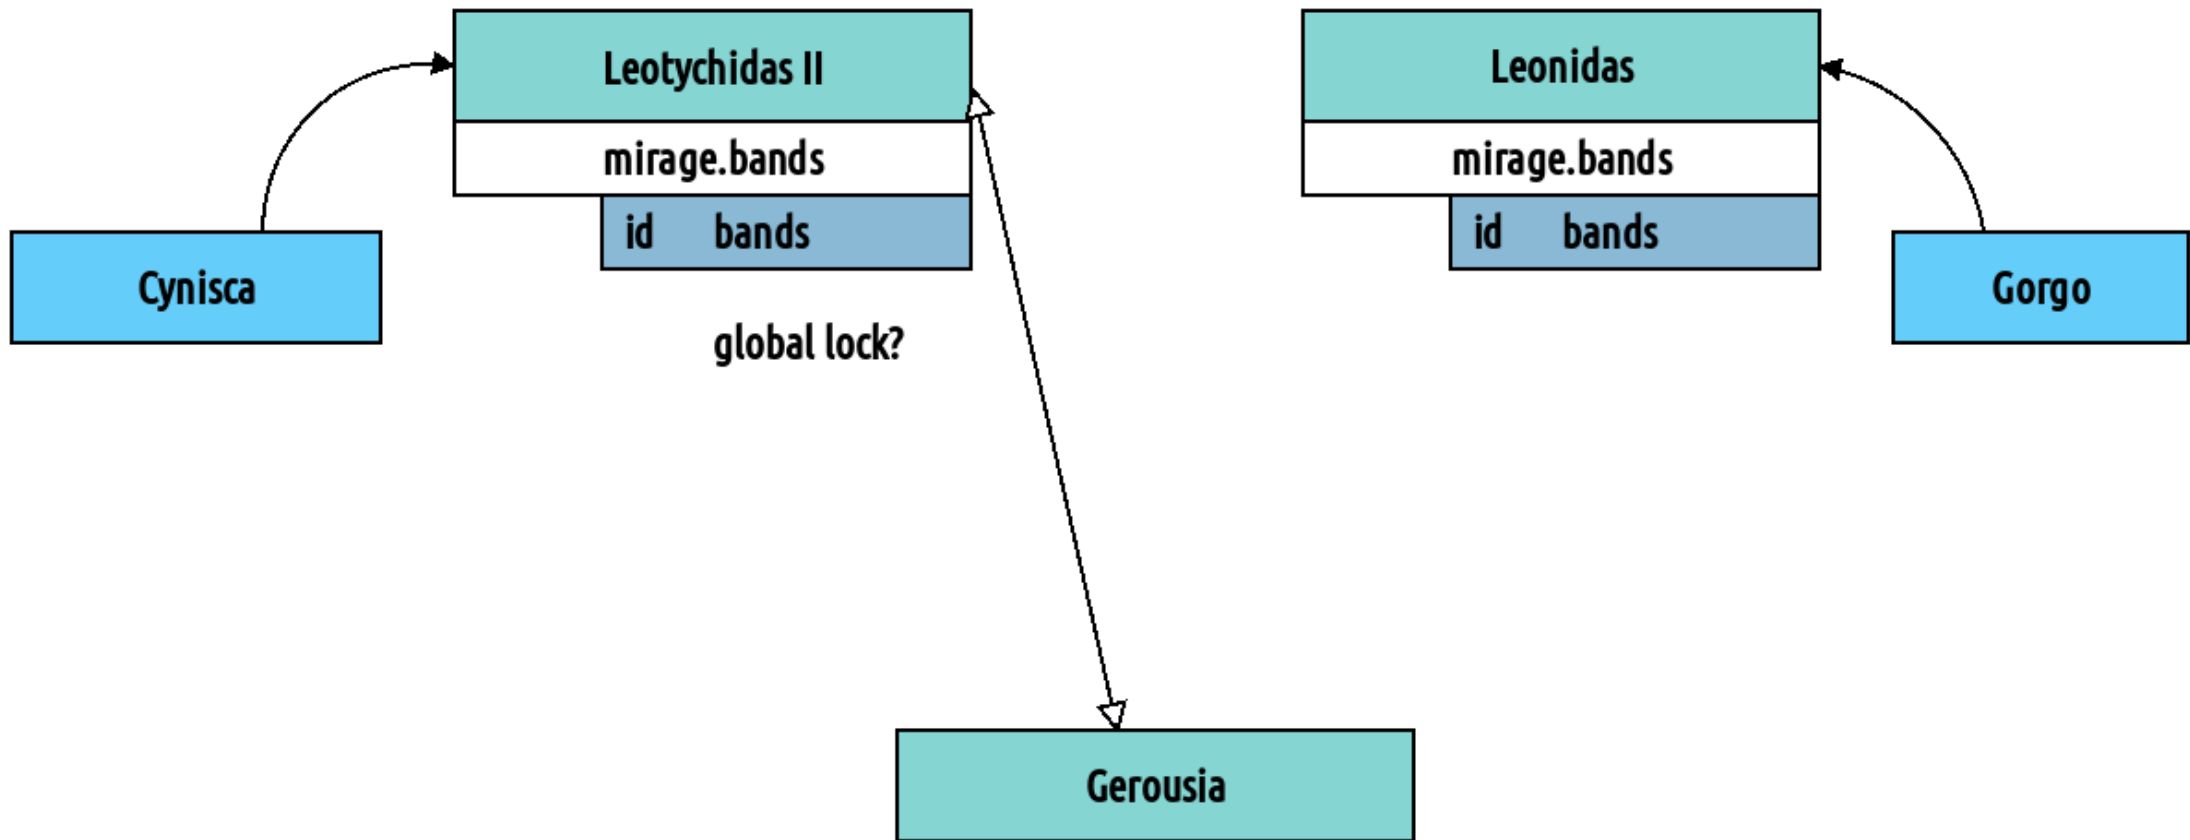


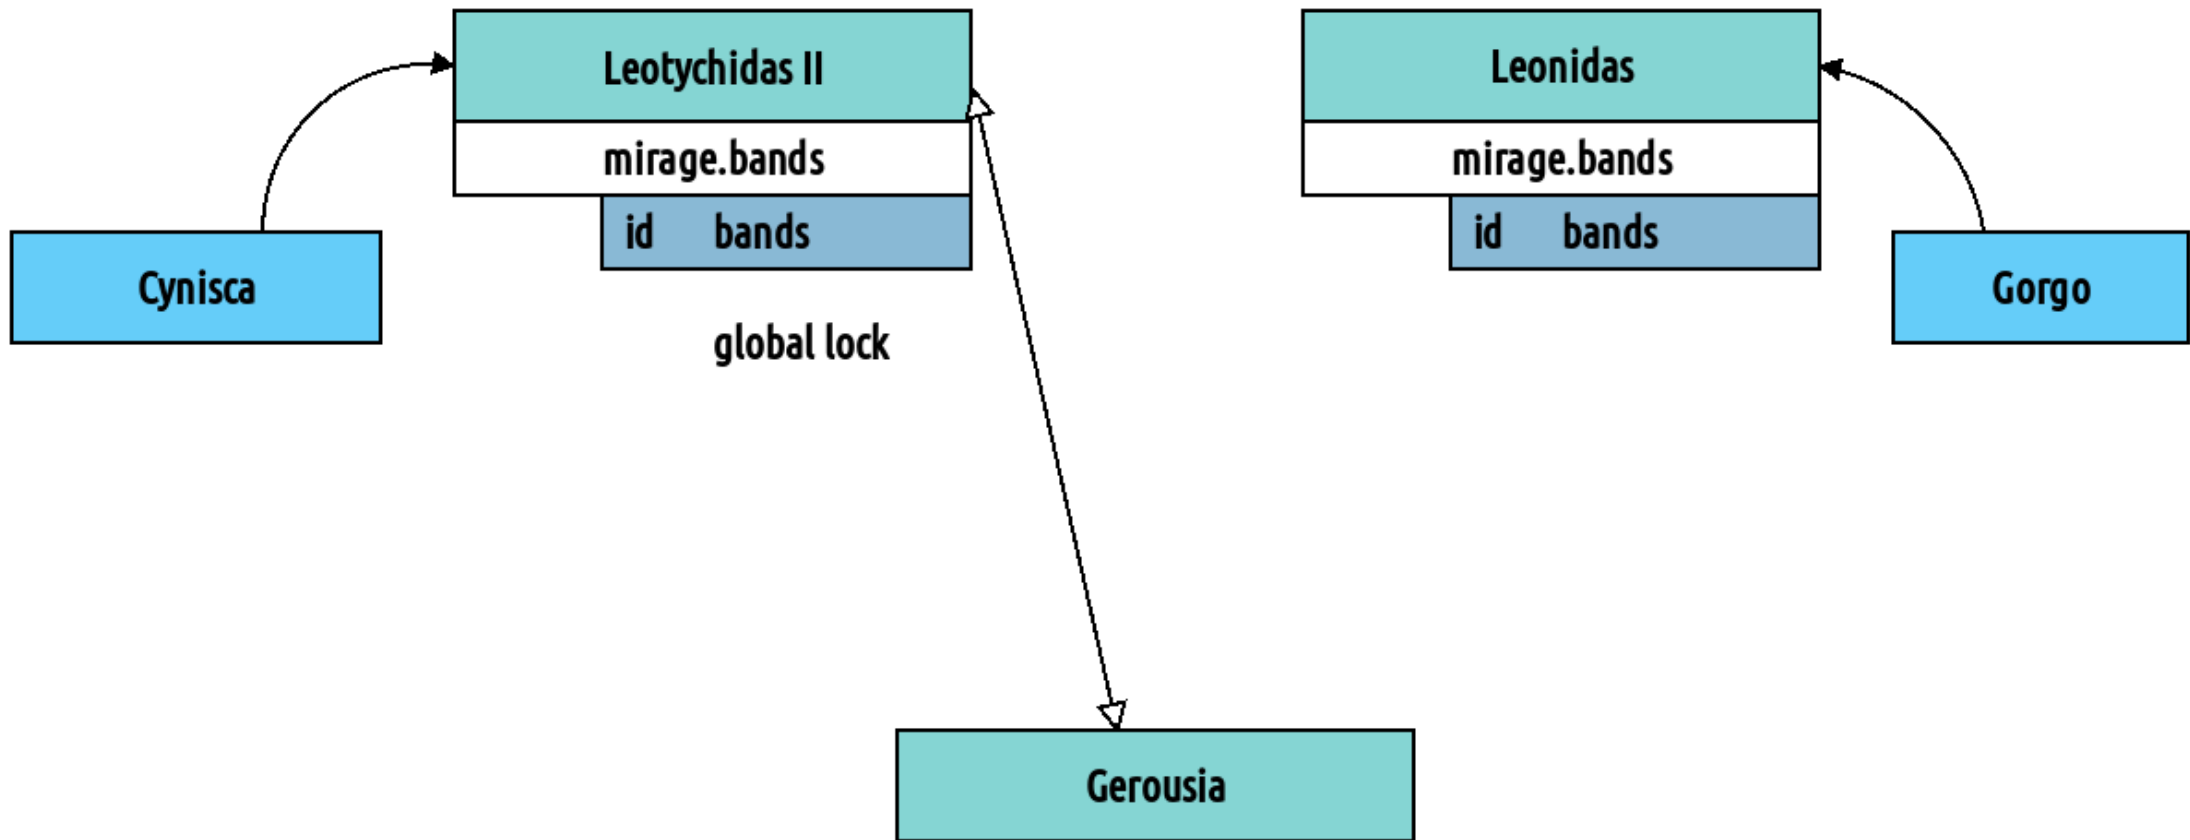


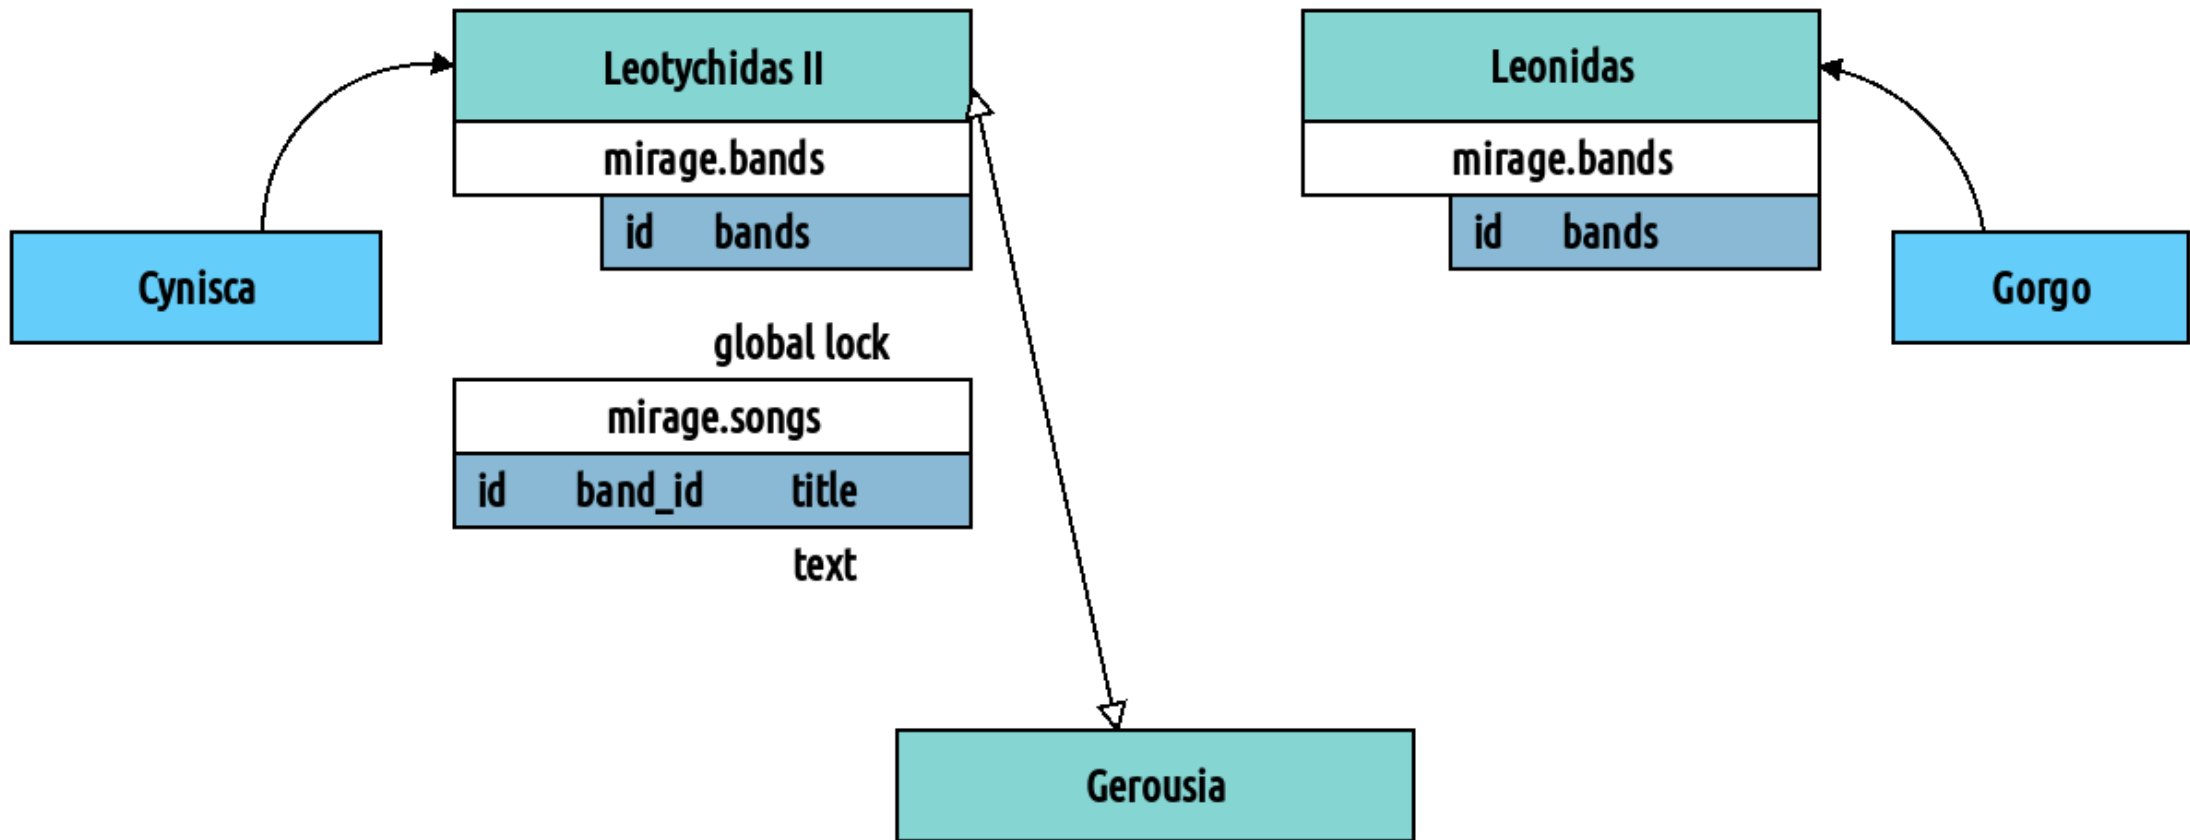


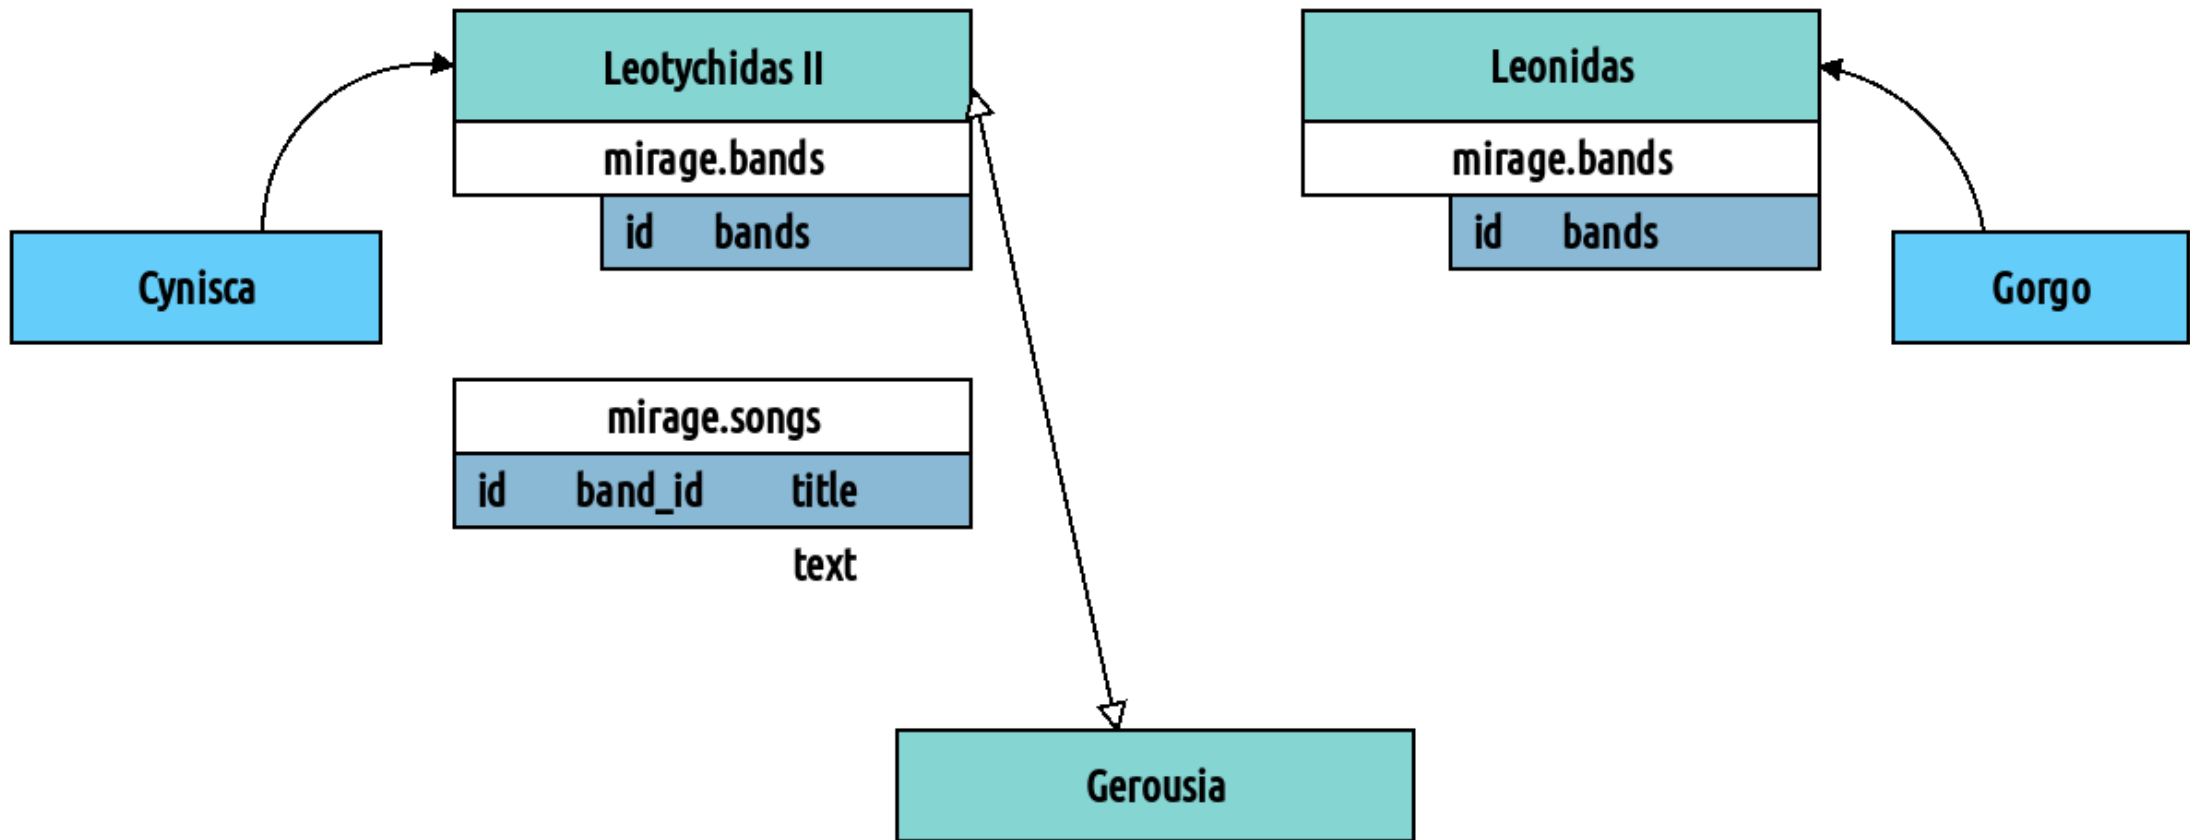


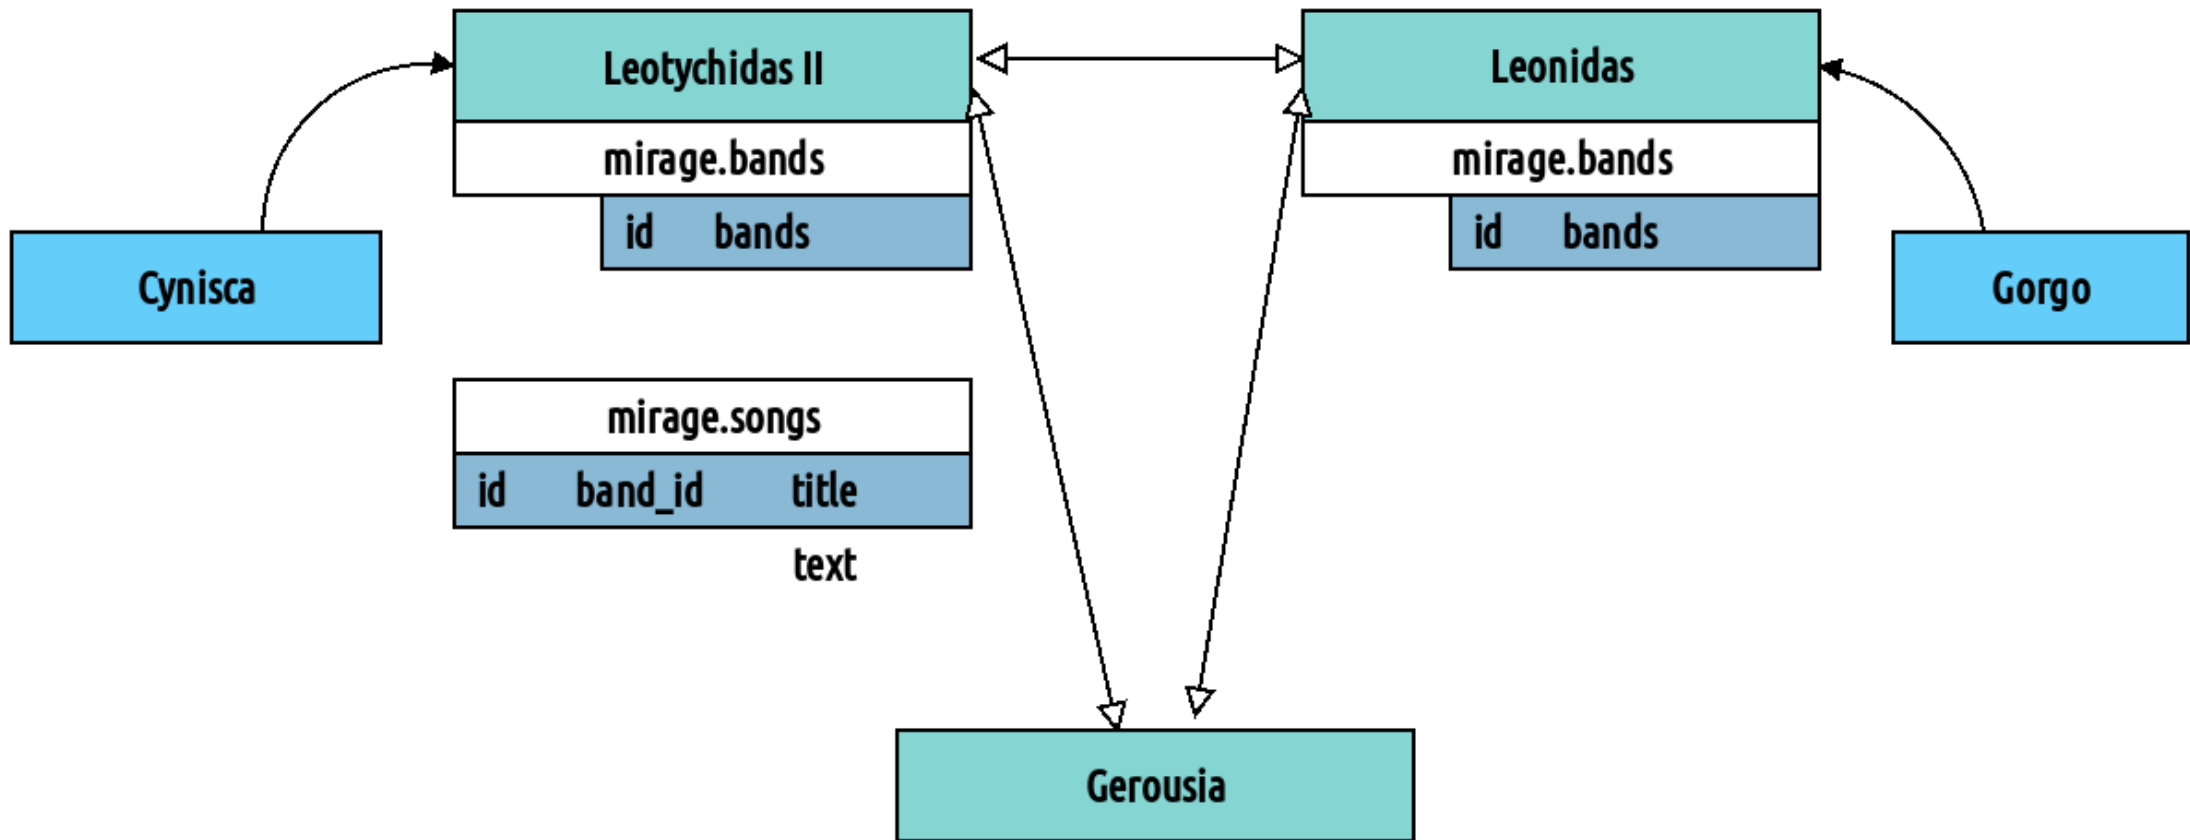


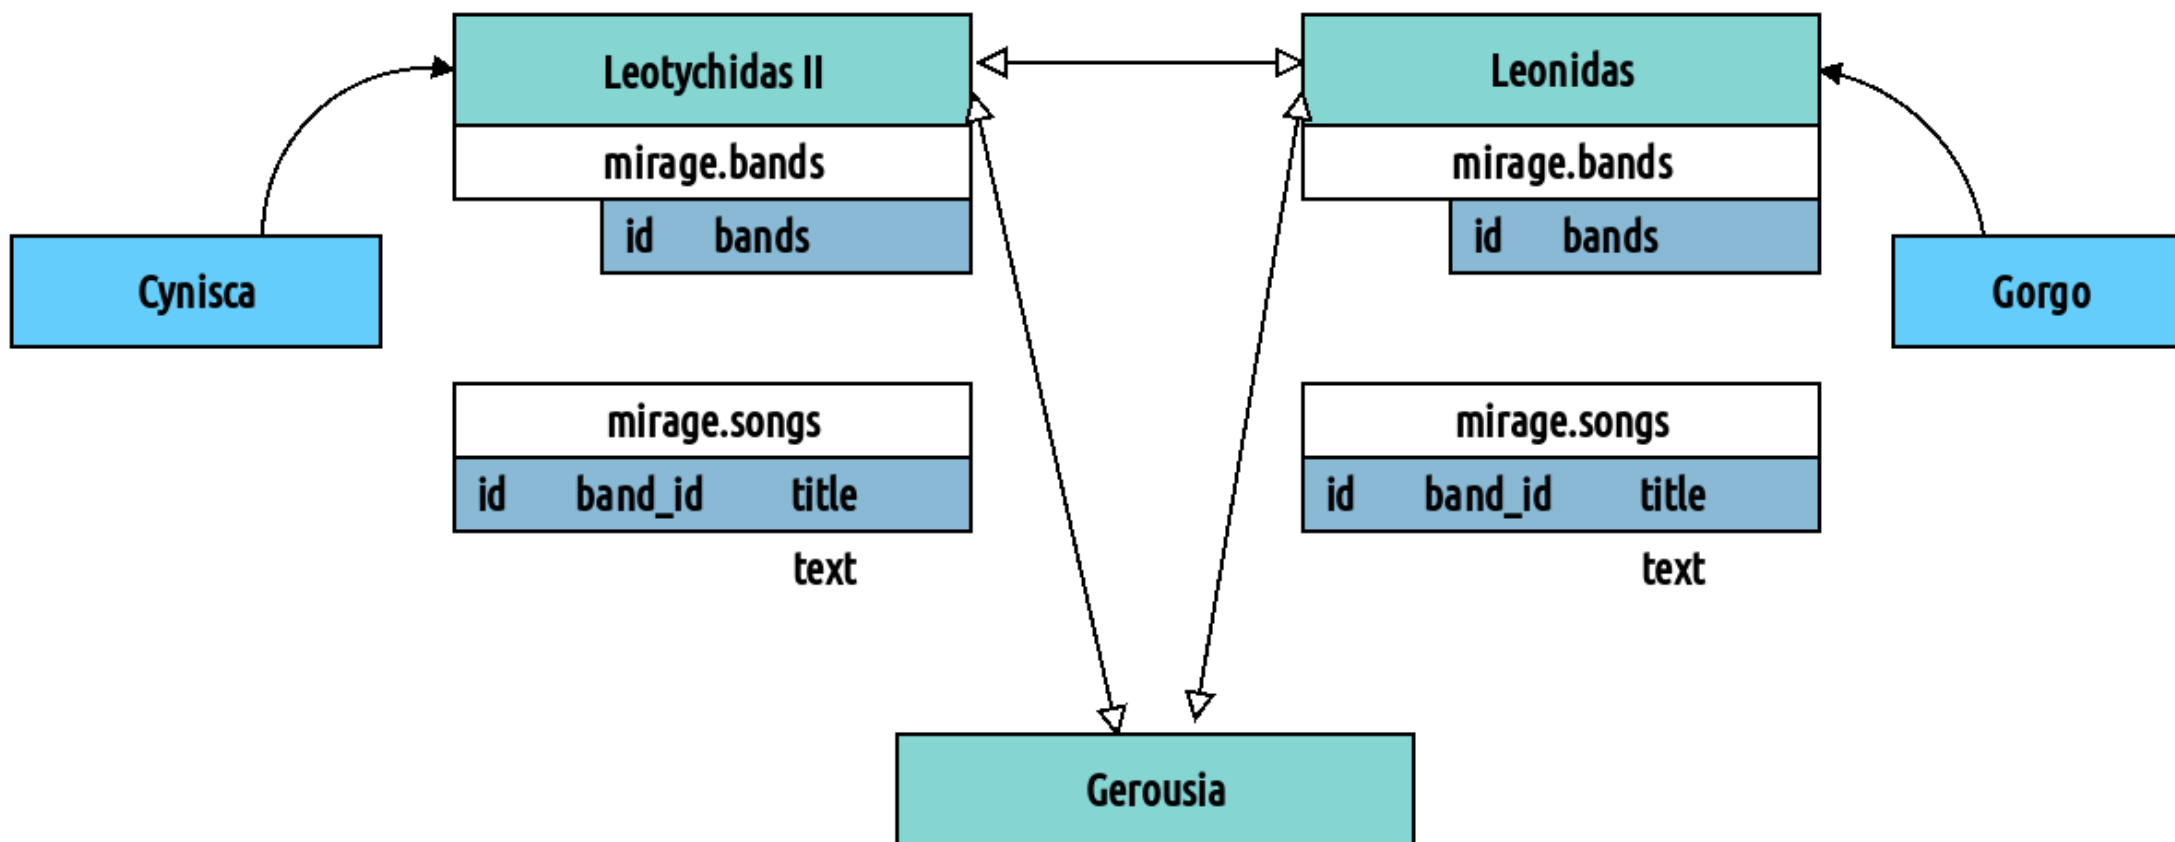






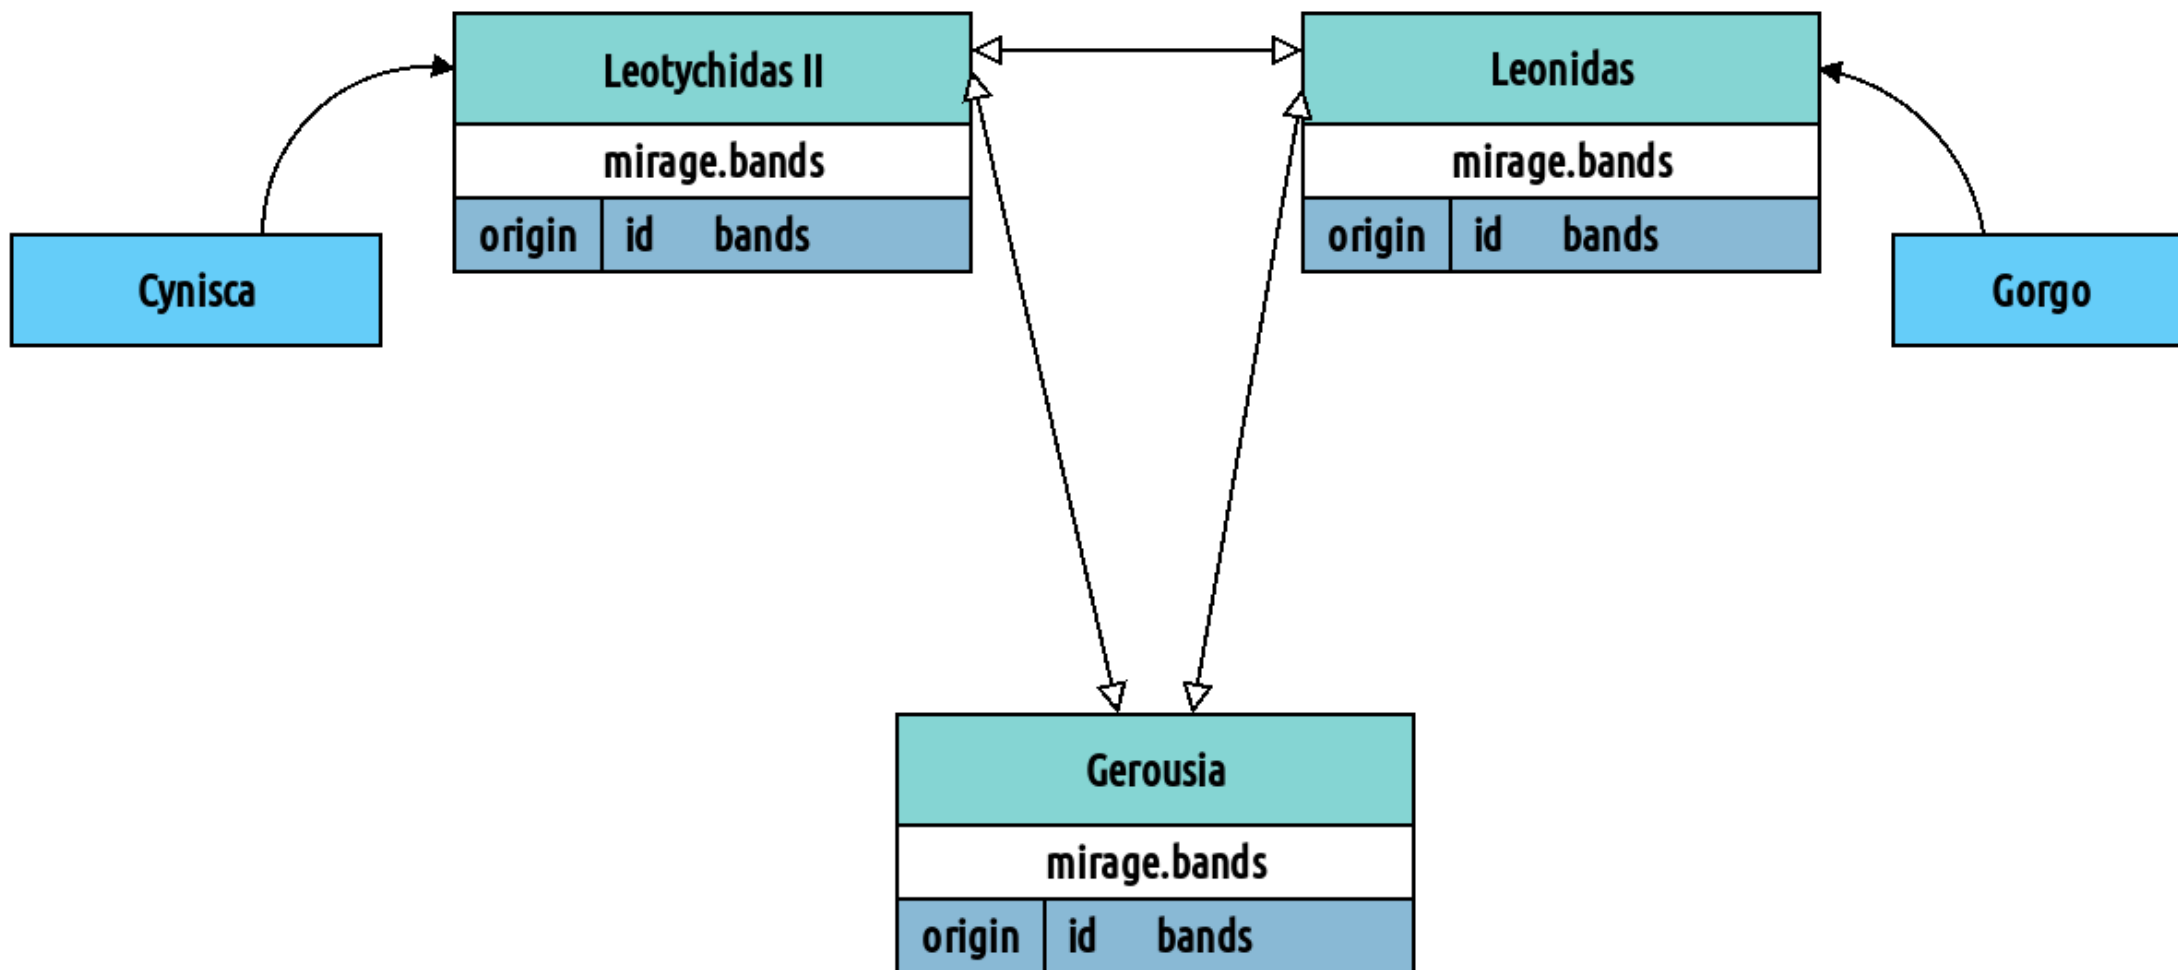




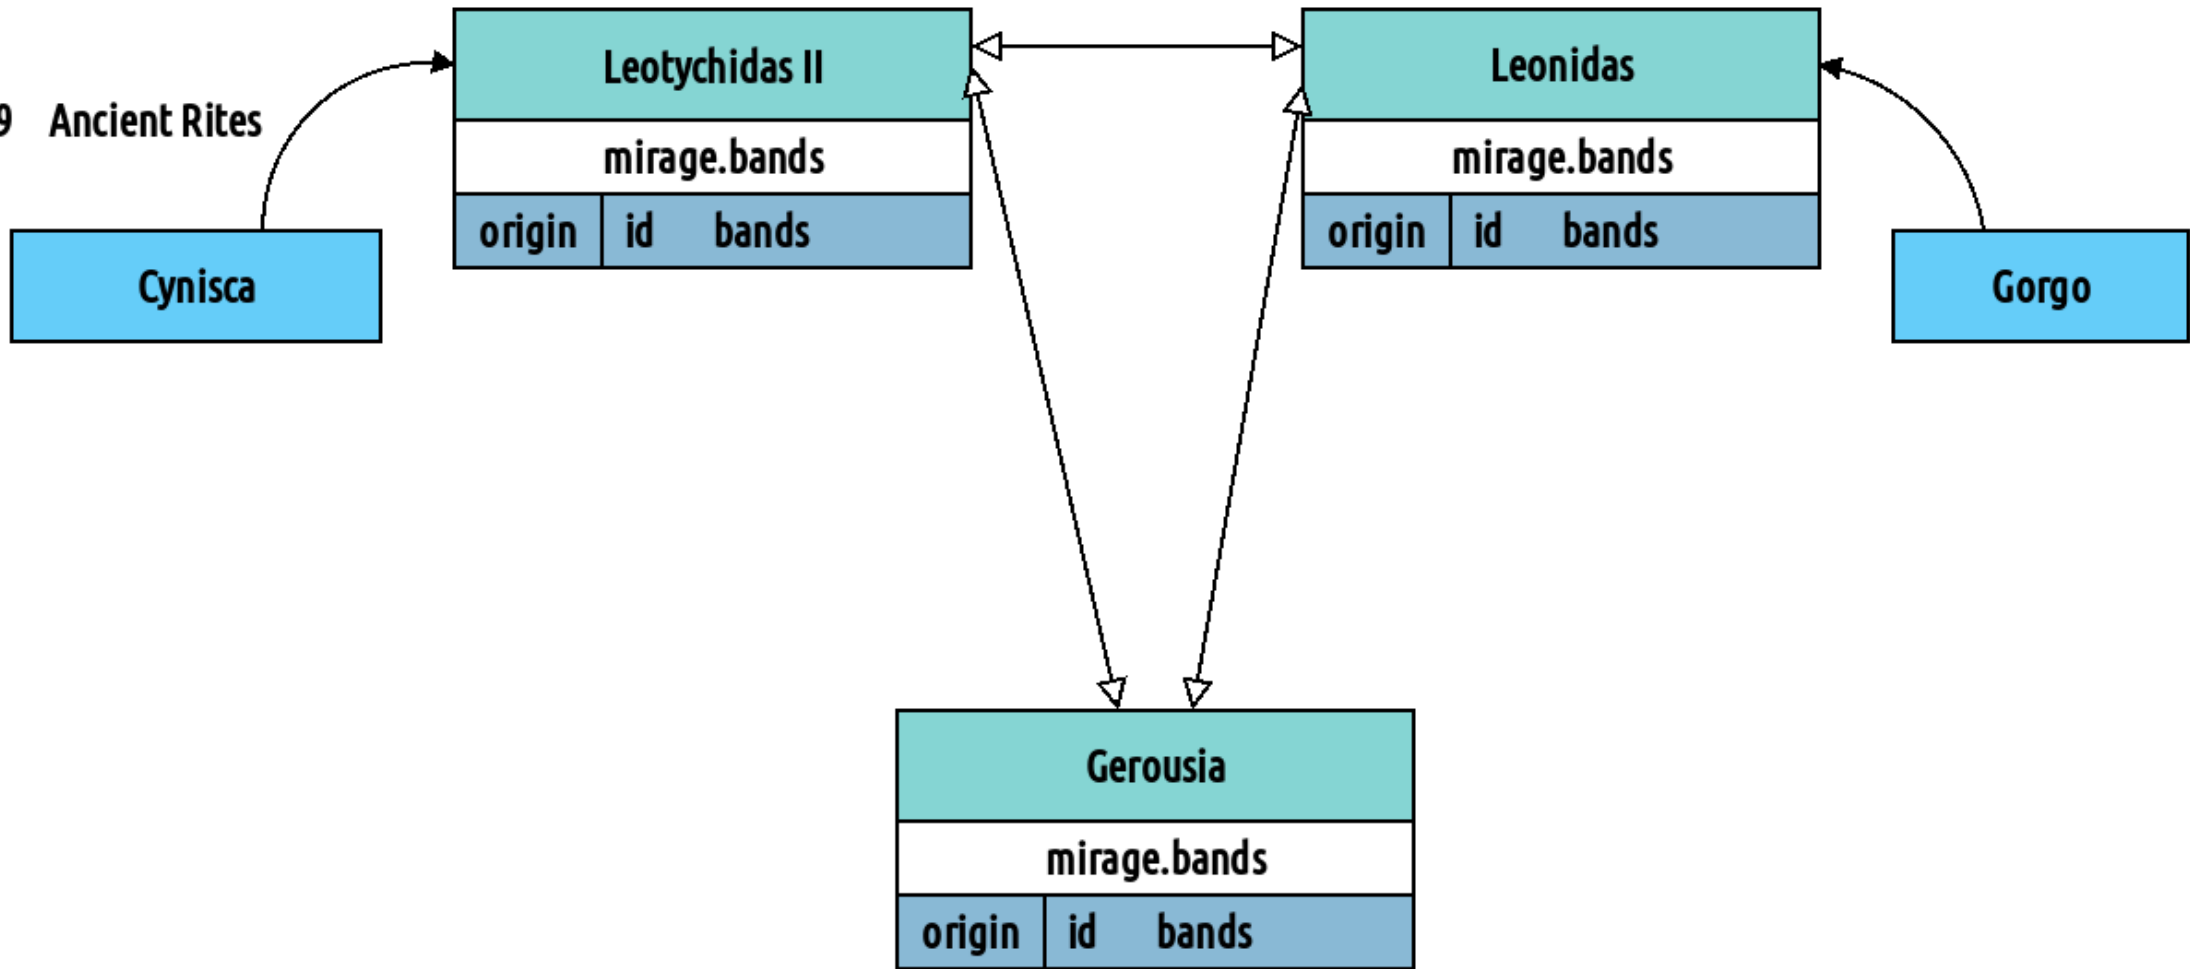


Deal with Conflicts using Consensus

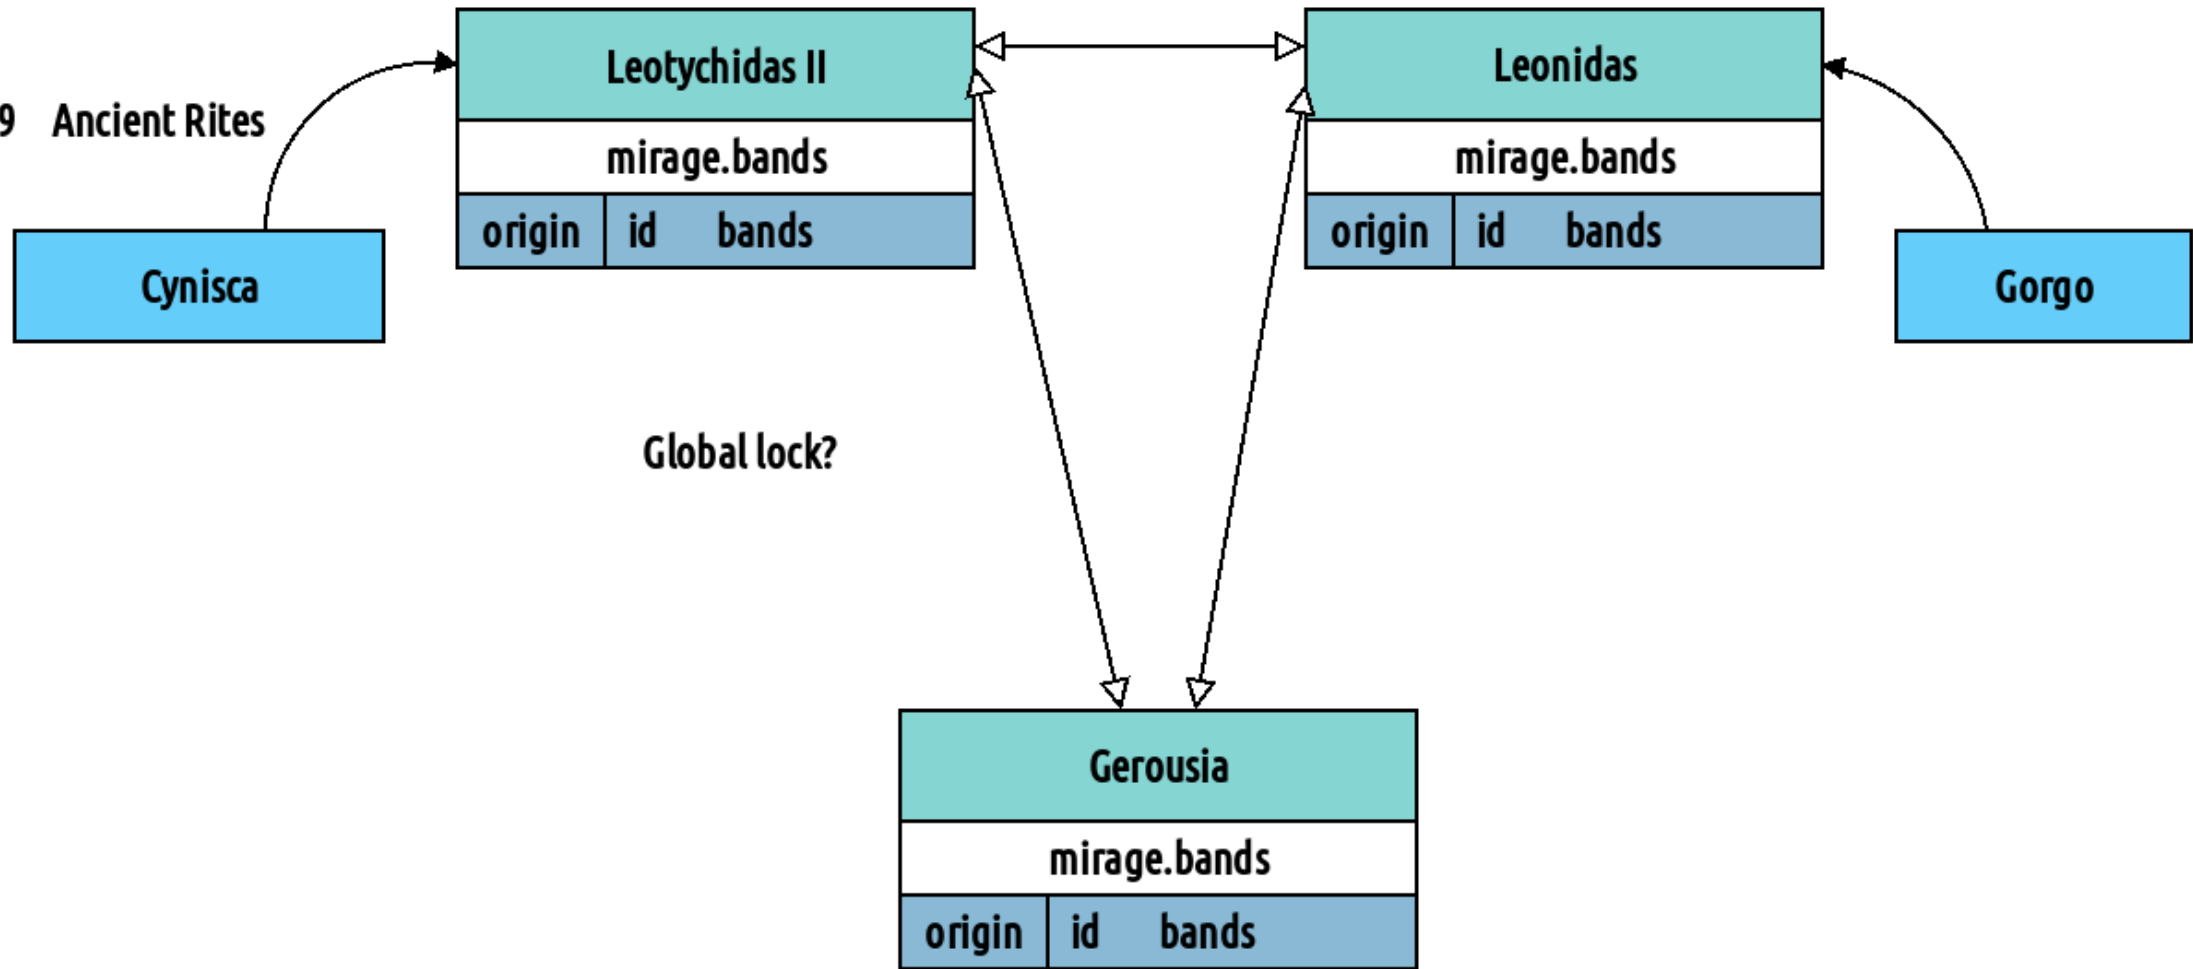




9 Ancient Rites





9 Ancient Rites



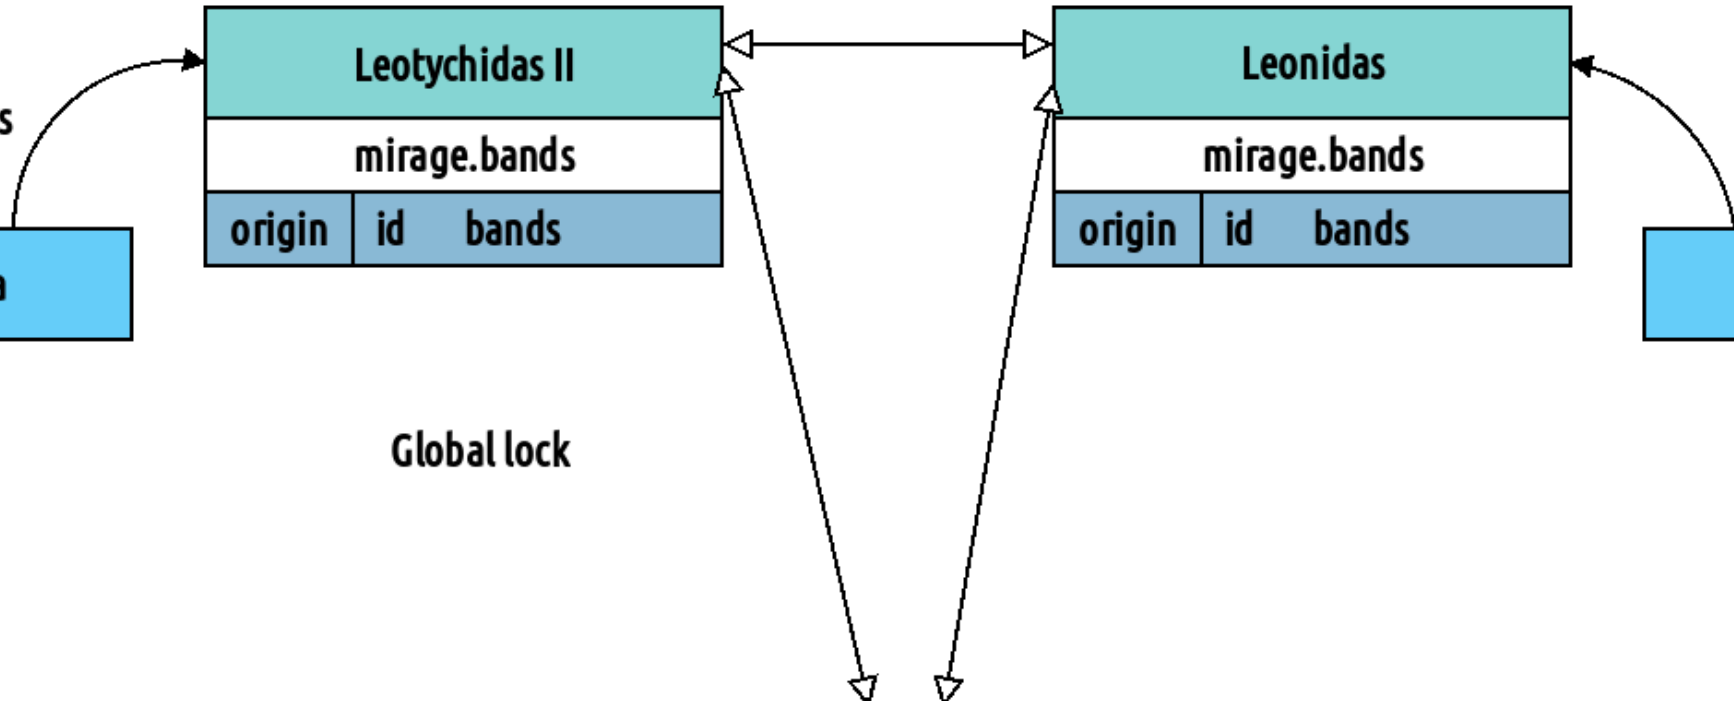

9 Ancient Rites

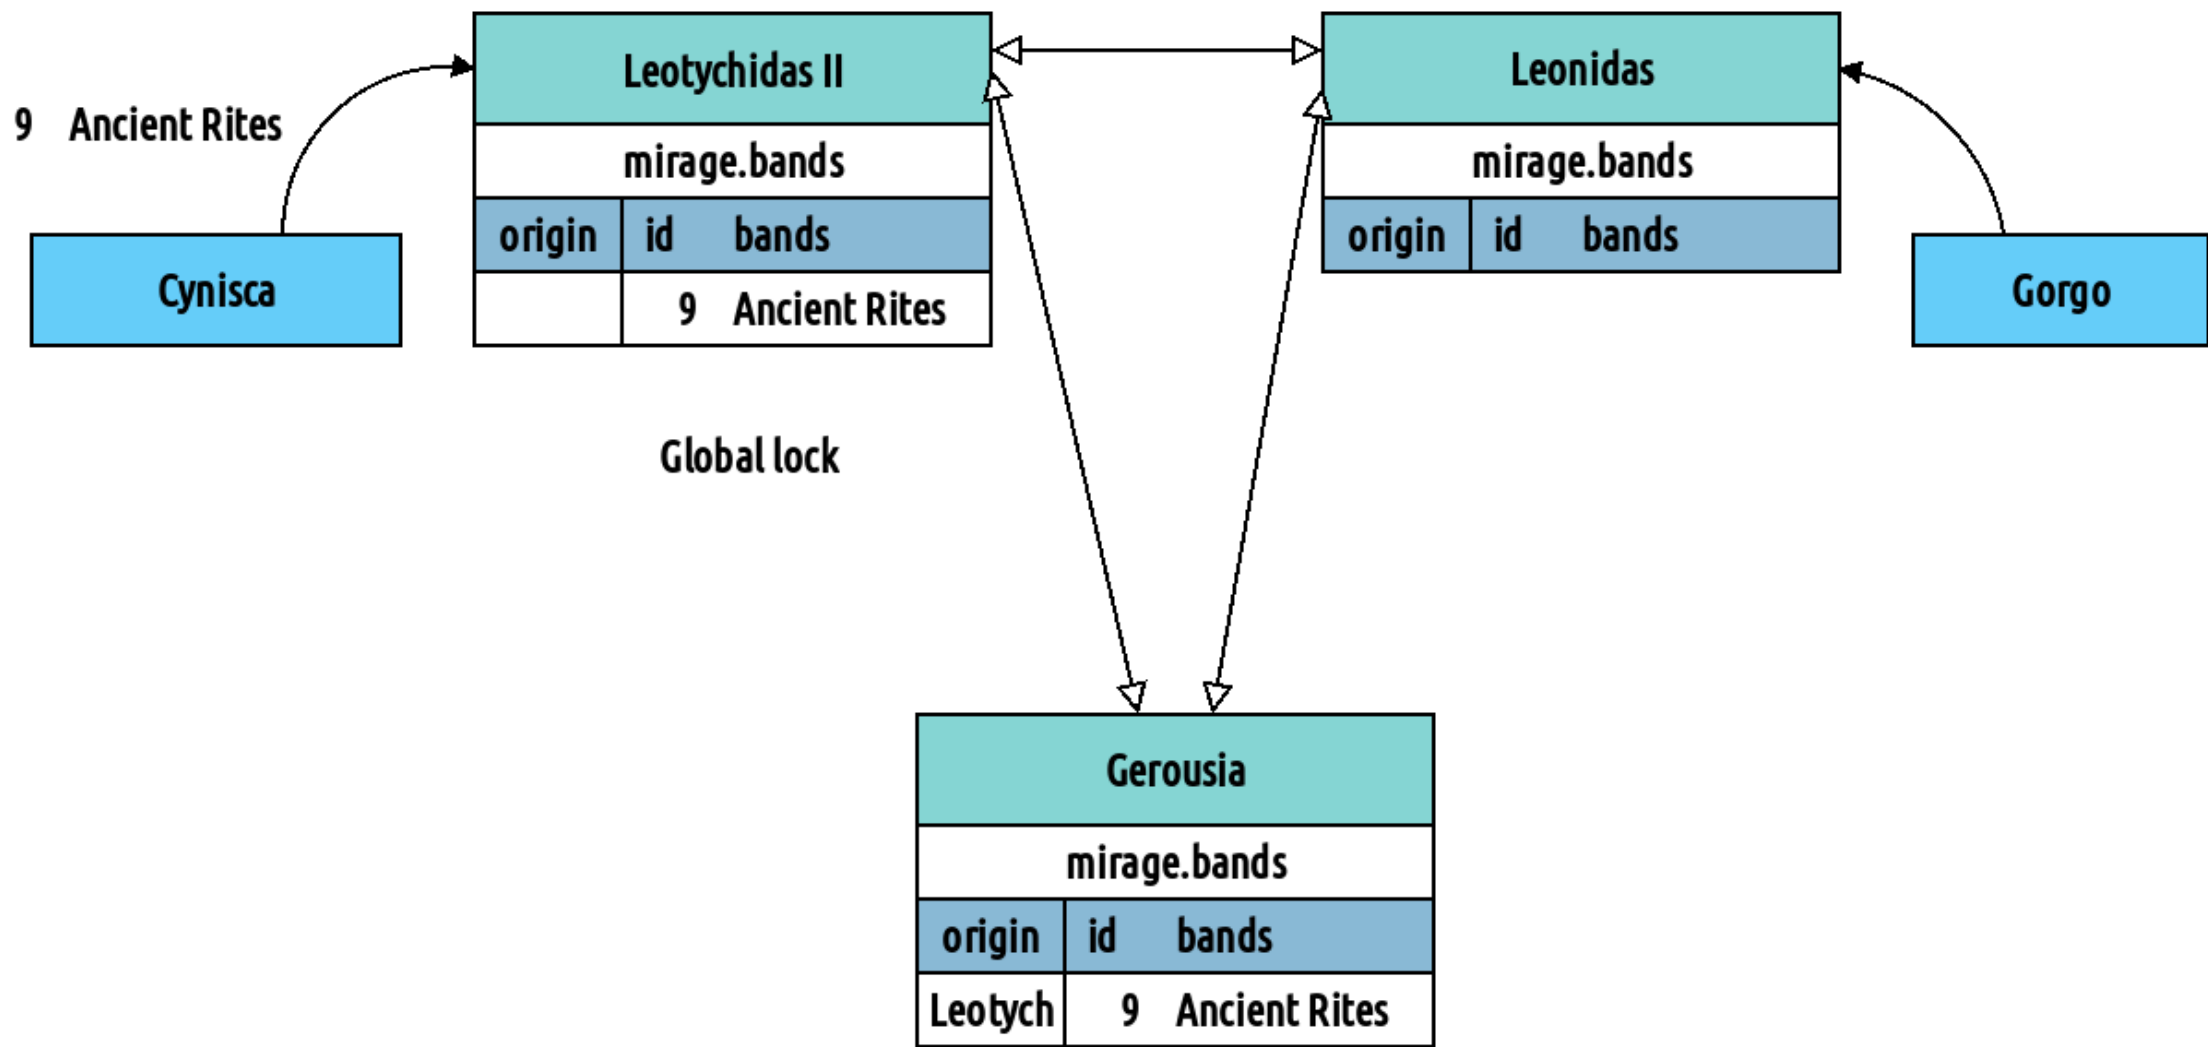
Cynisca




Gorgo

Global lock







Write in the majority

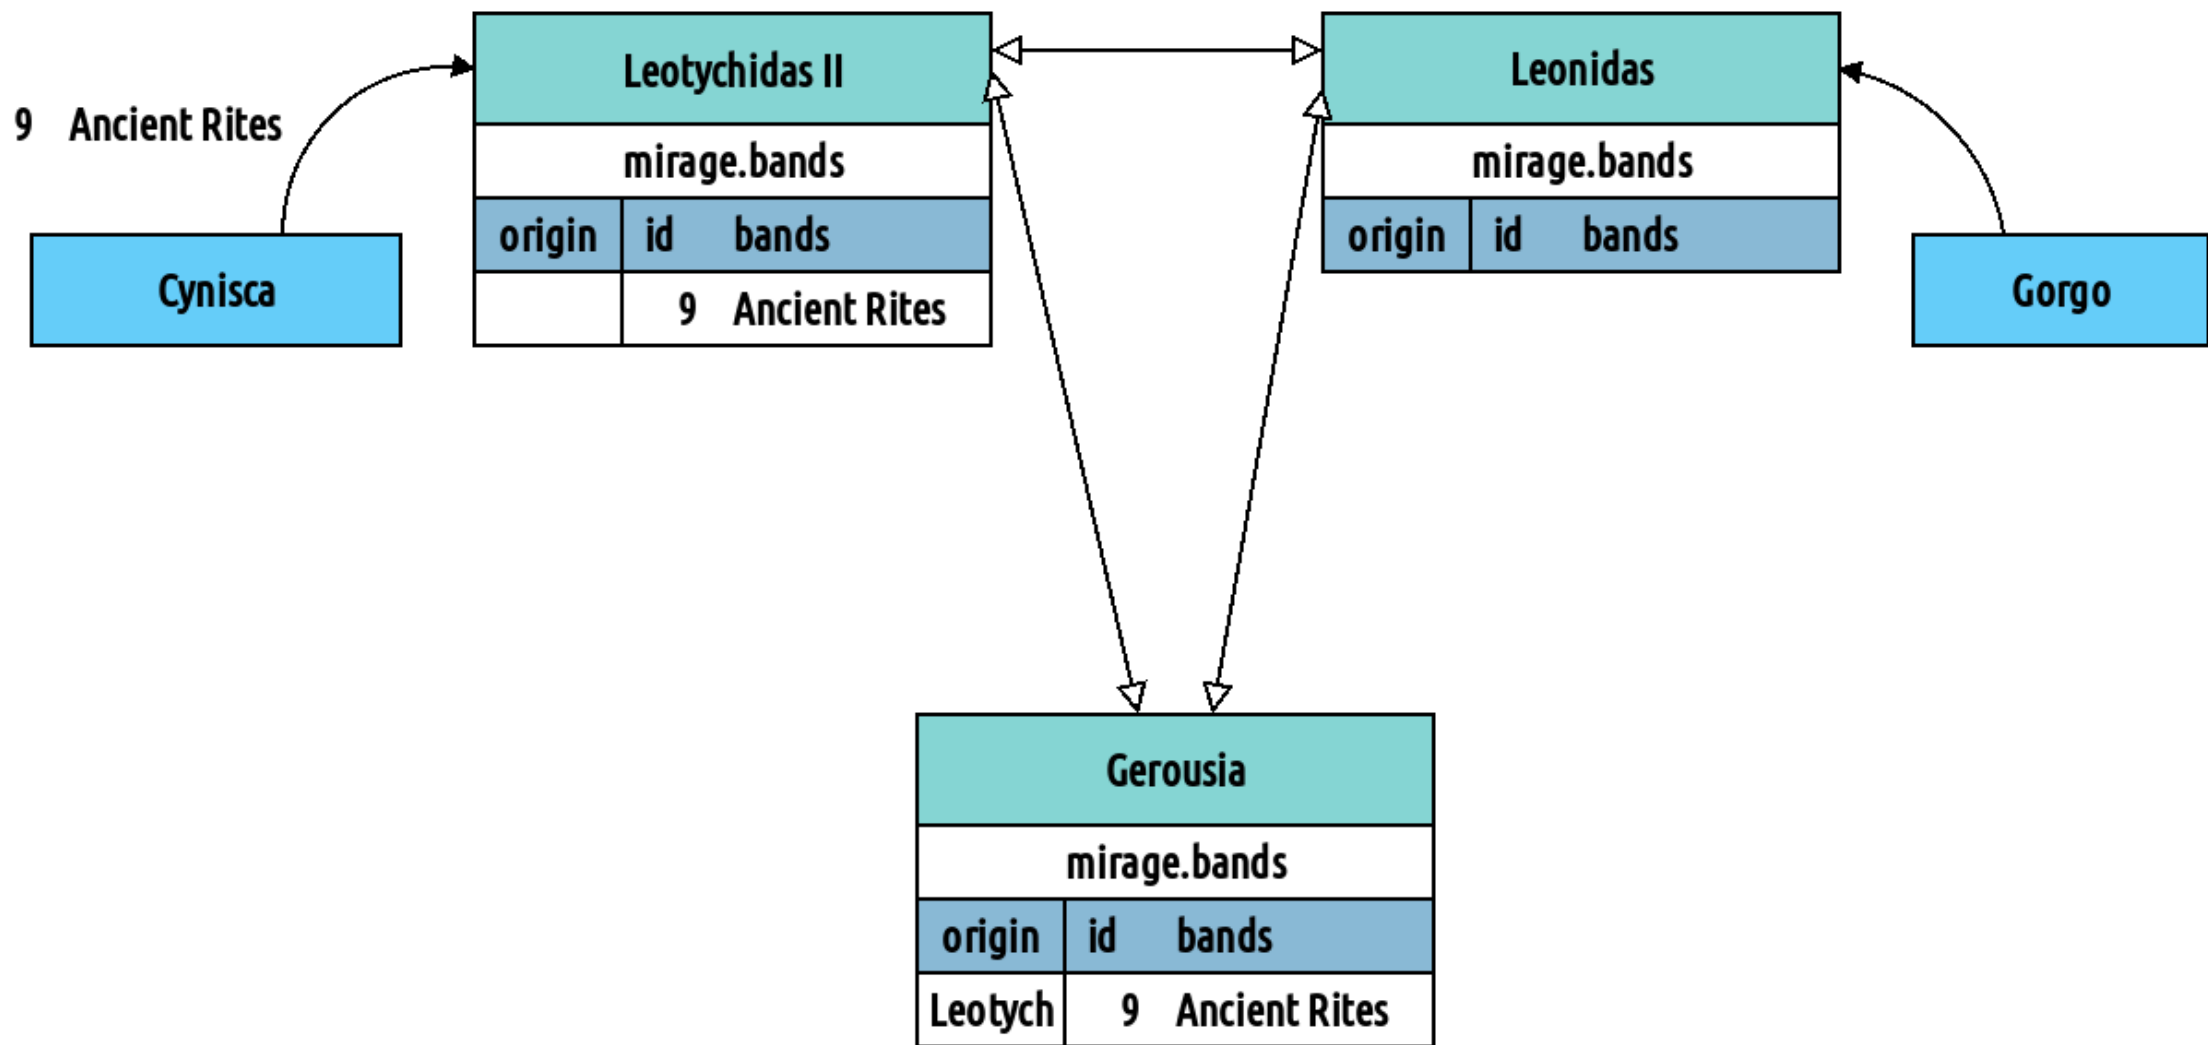


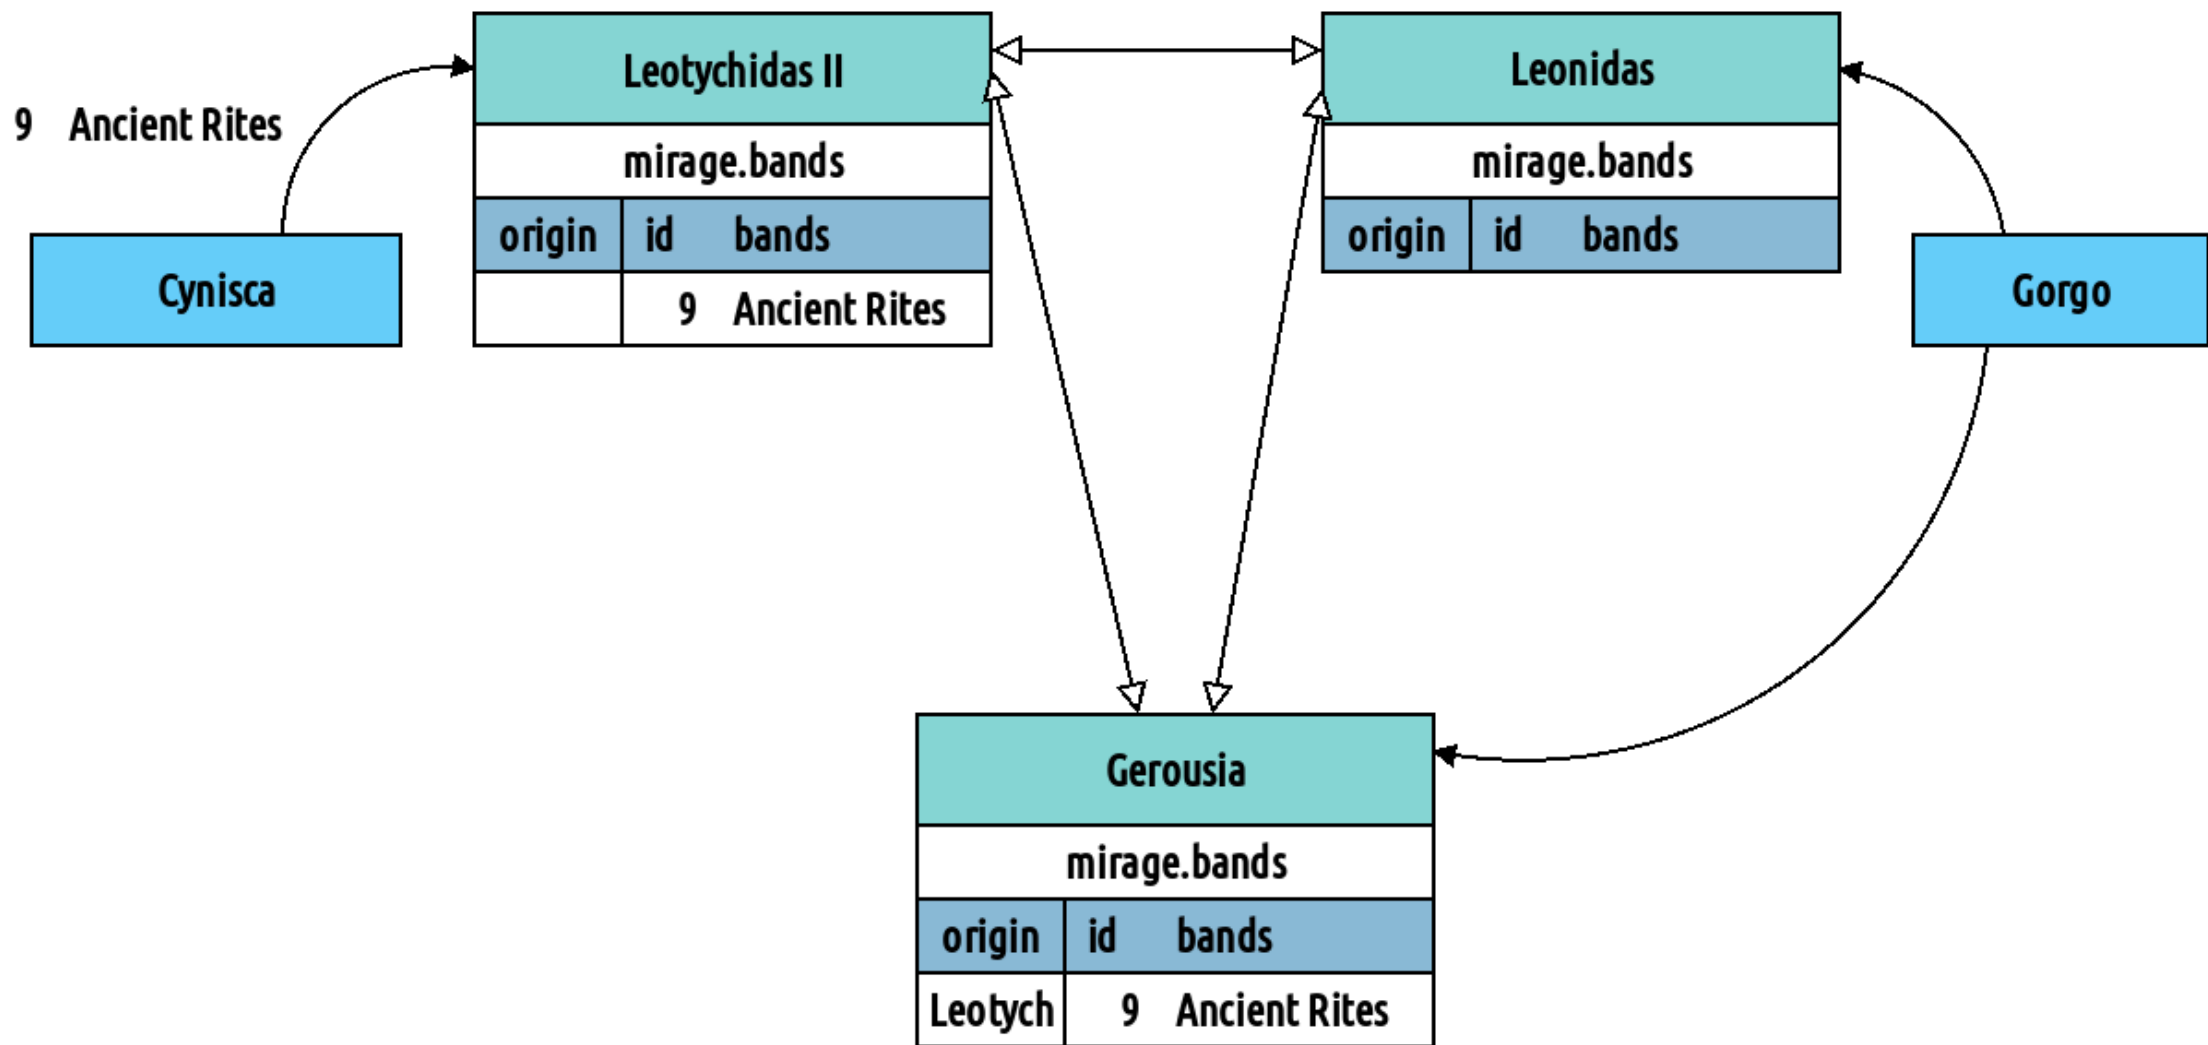
Write in the majority
Read from one node?



Write in the majority
Read from one node?
Read from majority?







9 Ancient Rites

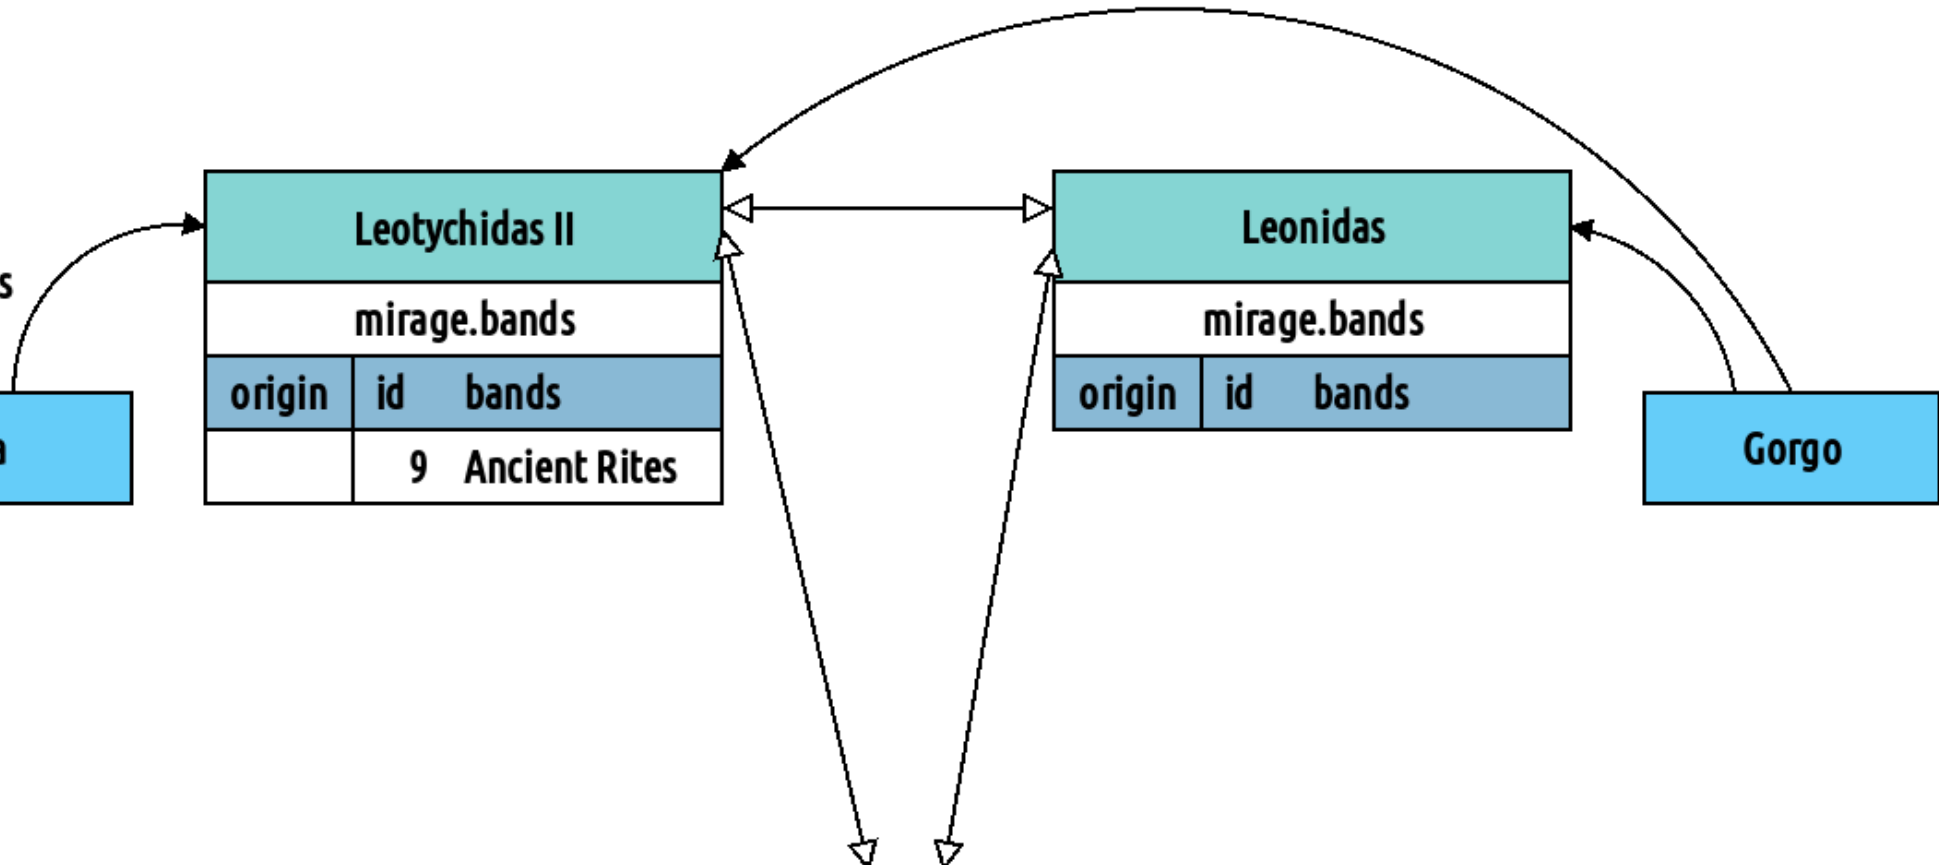
Cynisca

Leotychidas II		
mirage.bands		
origin	id	bands
	9	Ancient Rites

Leonidas		
mirage.bands		
origin	id	bands

Gorgo

Gerousia		
mirage.bands		
origin	id	bands
Leotych	9	Ancient Rites



9 Ancient Rites

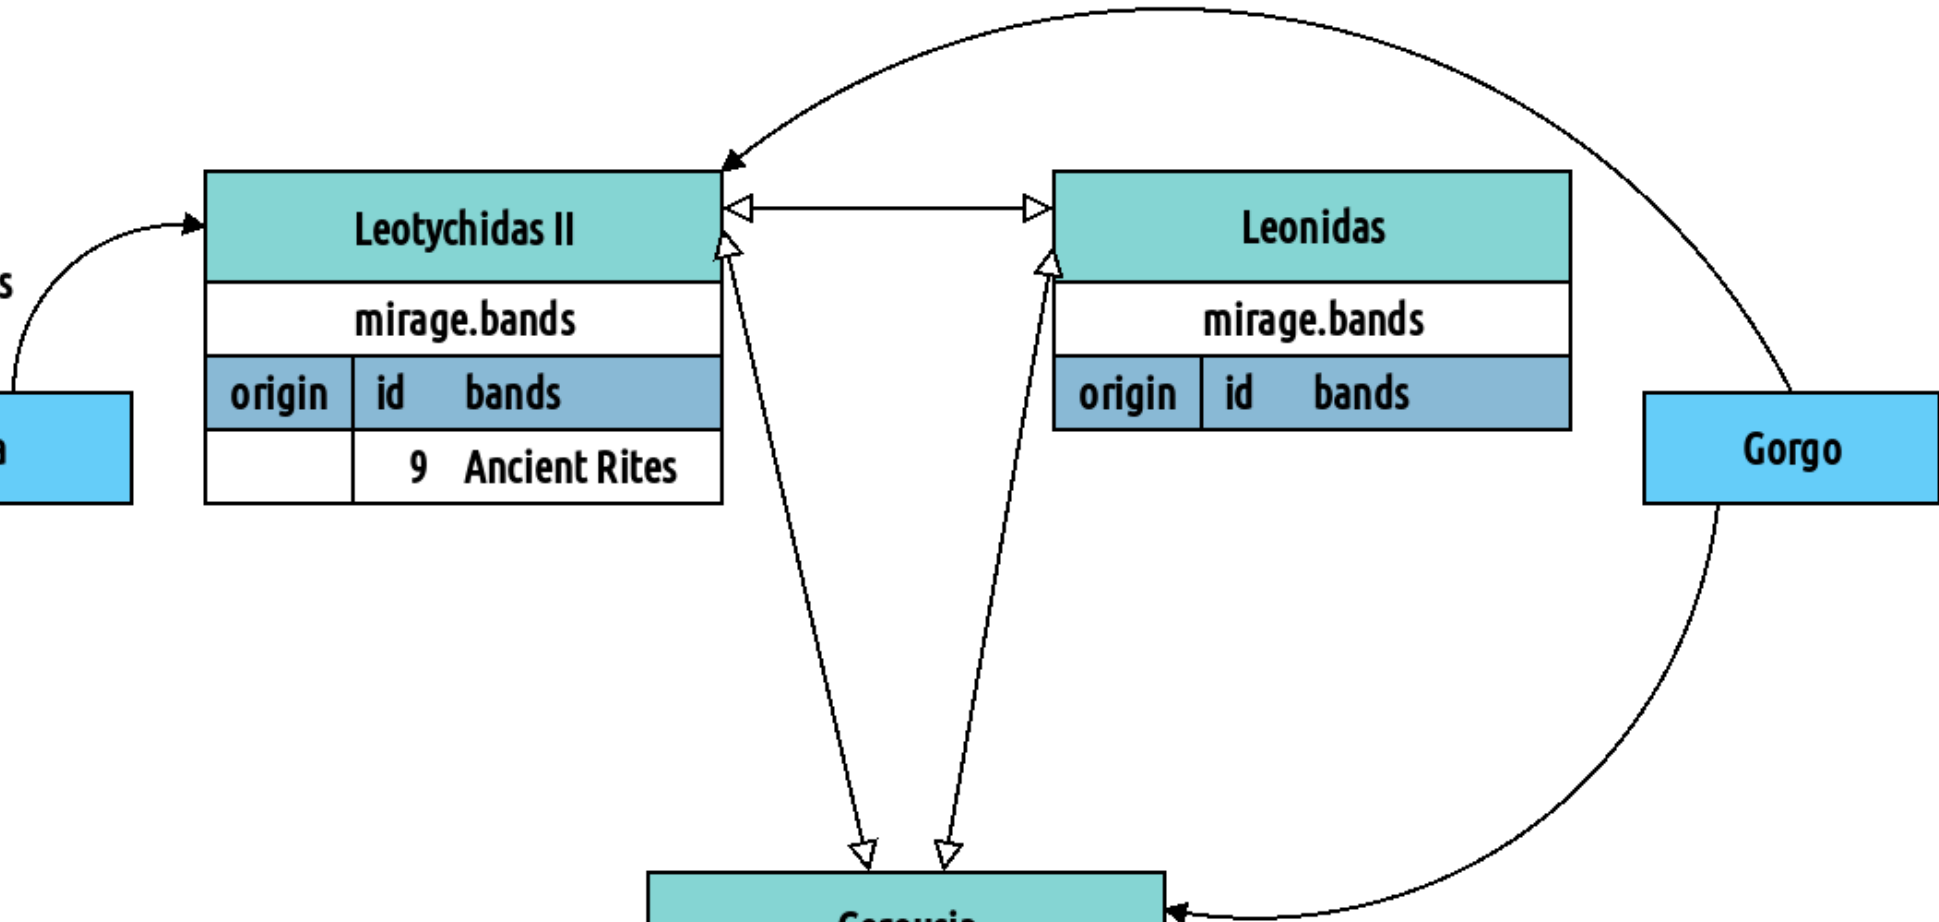
Cynisca

Leotychidas II		
mirage.bands		
origin	id	bands
	9	Ancient Rites

Leonidas		
mirage.bands		
origin	id	bands

Gorgo

Gerousia		
mirage.bands		
origin	id	bands
Leotych	9	Ancient Rites



9 Ancient Rites

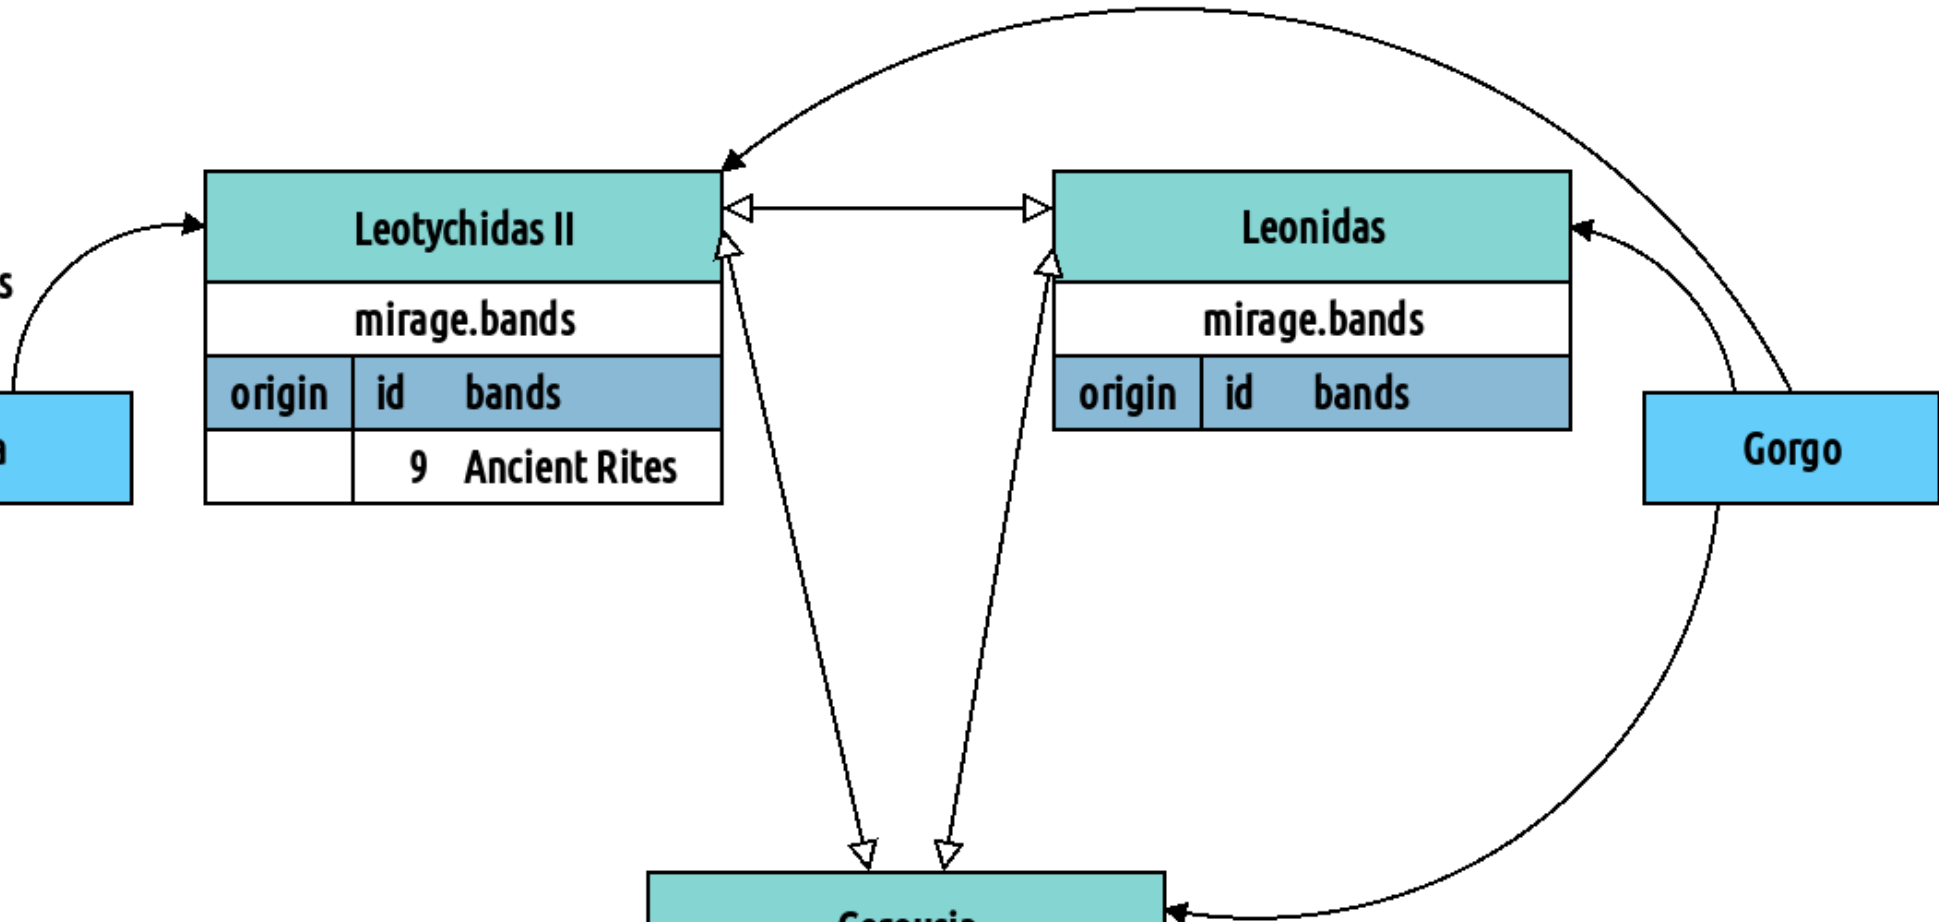
Cynisca

Leotychidas II		
mirage.bands		
origin	id	bands
	9	Ancient Rites

Leonidas		
mirage.bands		
origin	id	bands

Gorgo

Gerousia		
mirage.bands		
origin	id	bands
Leotych	9	Ancient Rites



9 Ancient Rites

Cynisca

Leotychidas II		
mirage.bands		
origin	id	bands
	9	Ancient Rites

Leonidas		
mirage.bands		
origin	id	bands

Gorgo

Gerousia		
mirage.bands		
origin	id	bands
Leotych	9	Ancient Rites

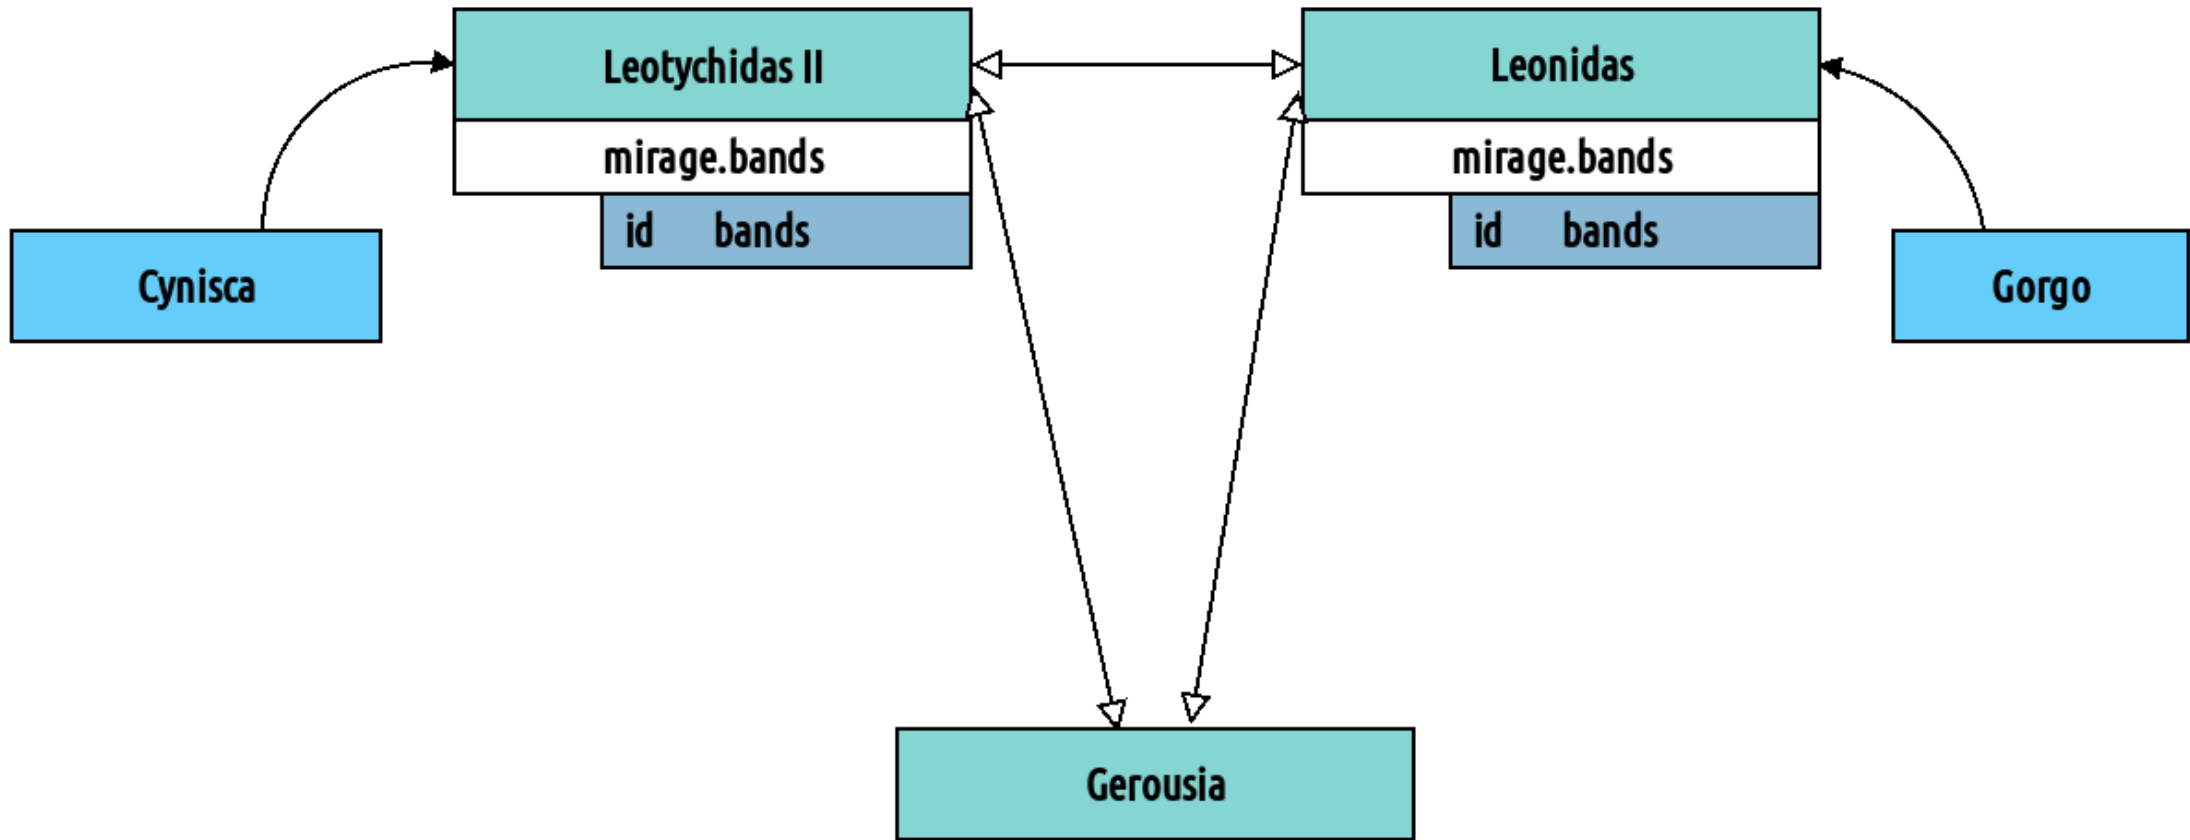
Band names curated by Jeremy Swist
The Spartan Metal Mirage

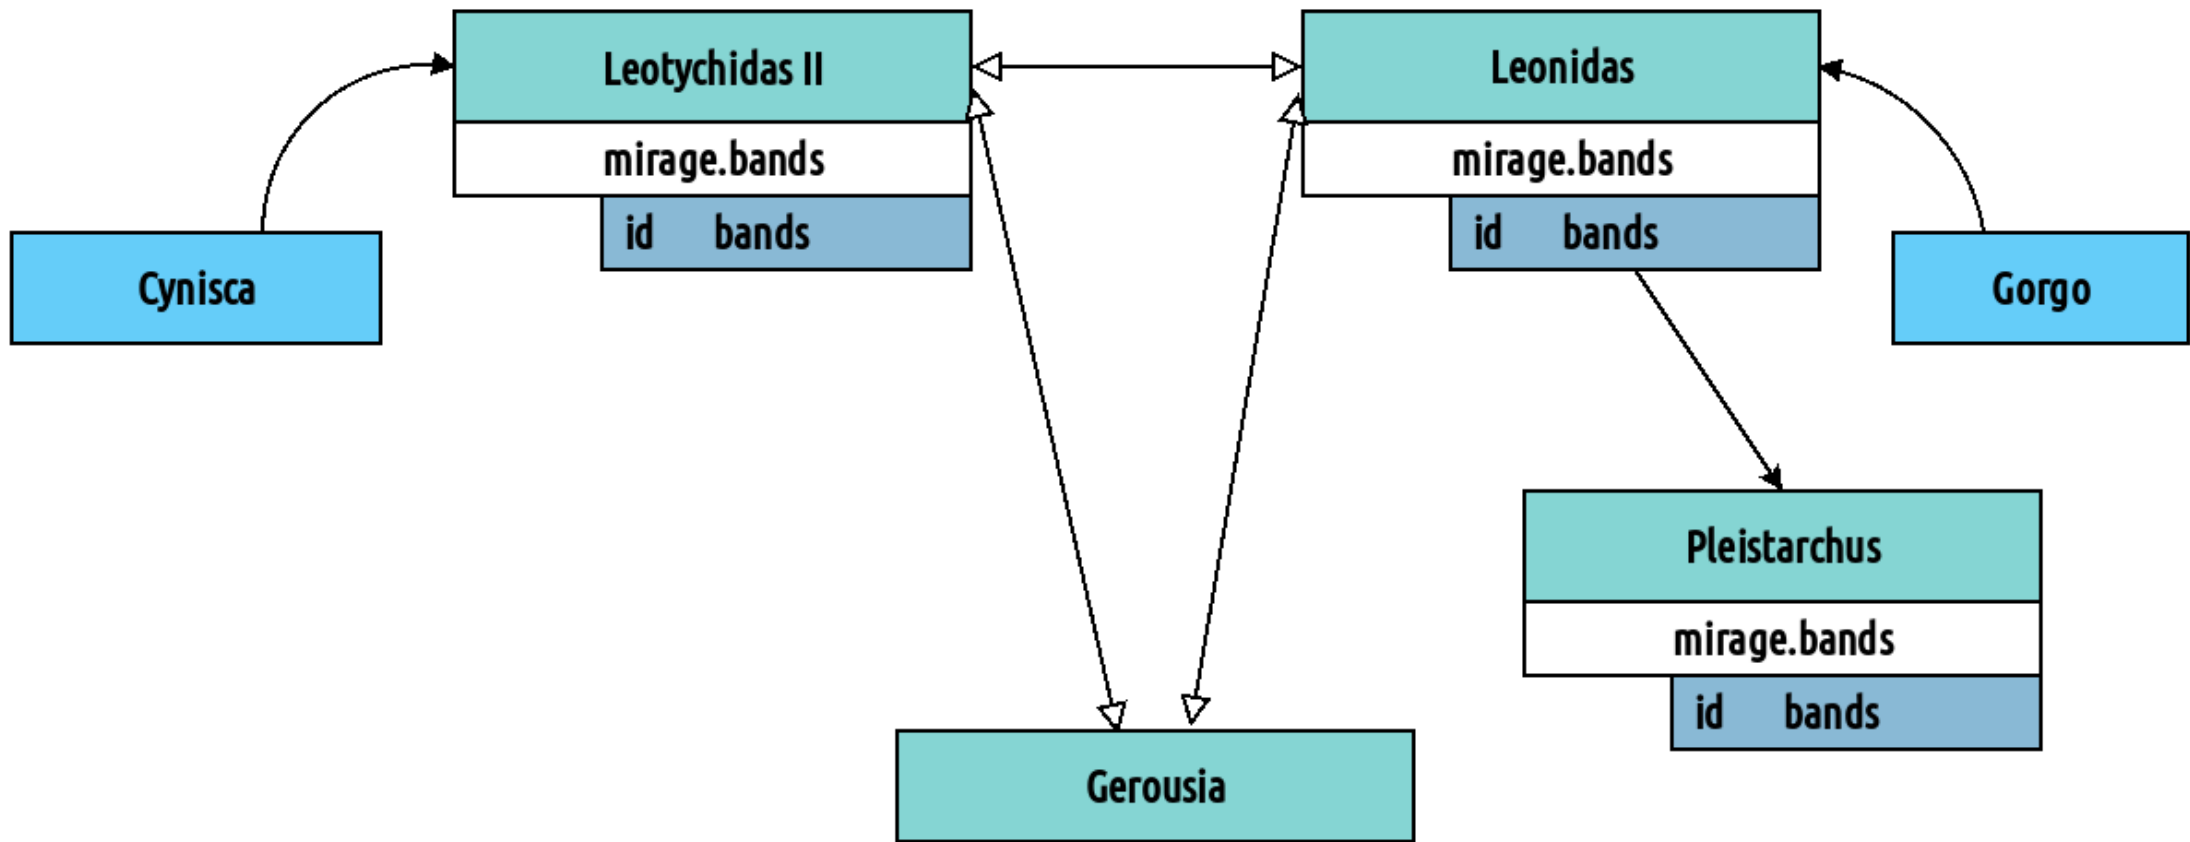
300

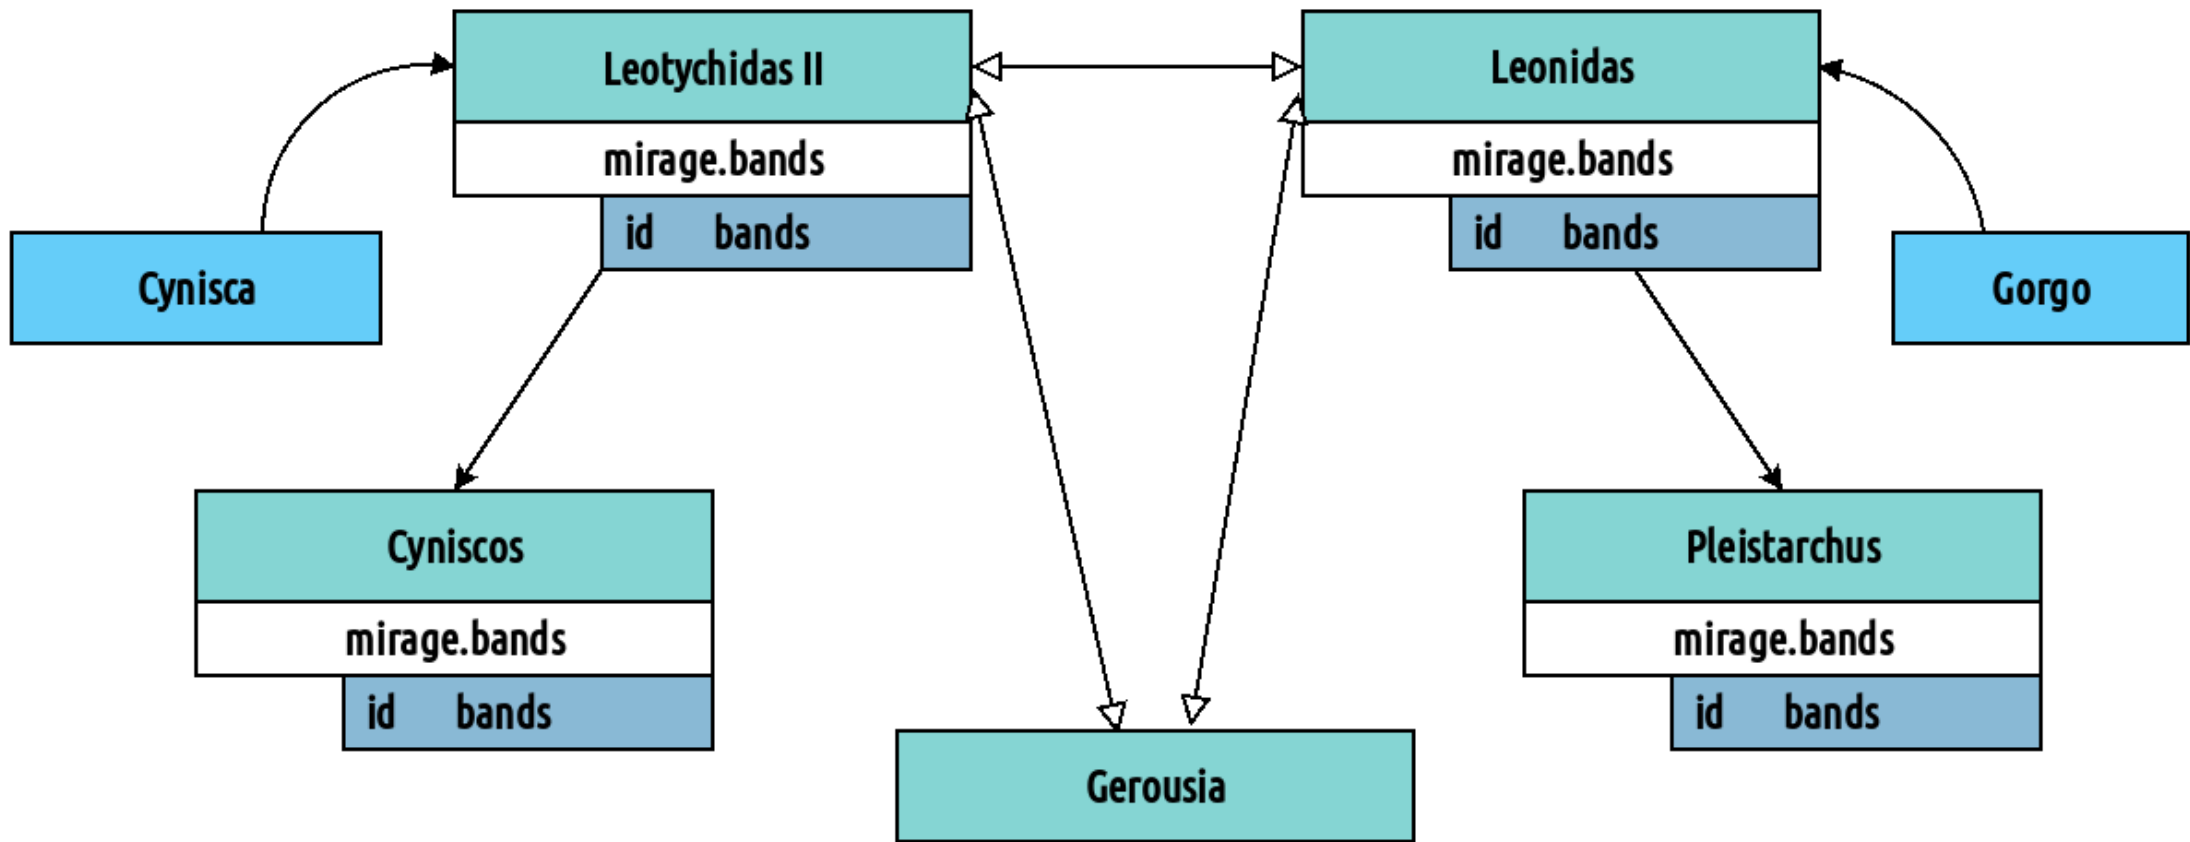


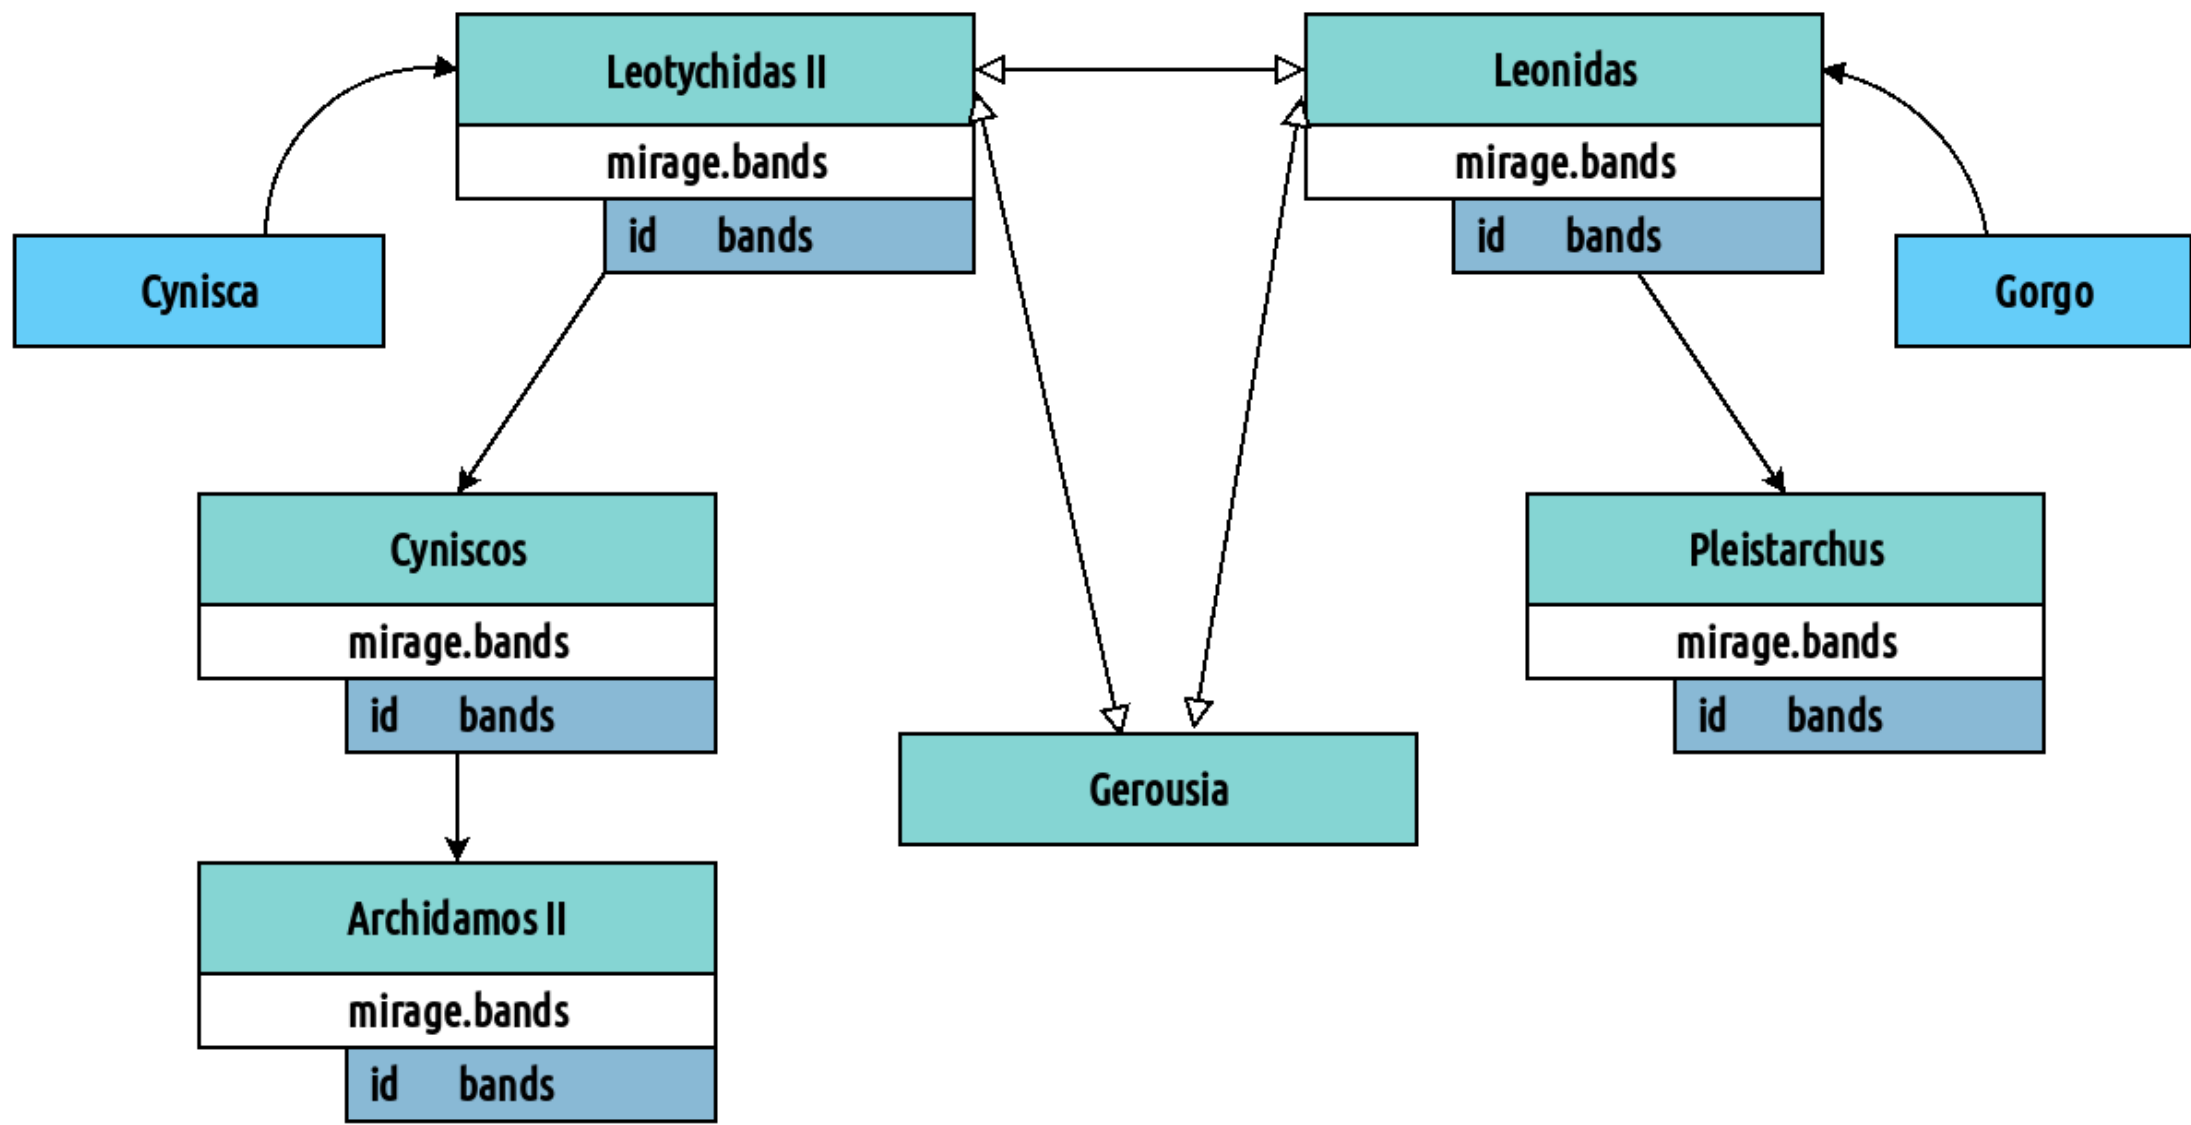
Disaster Recovery

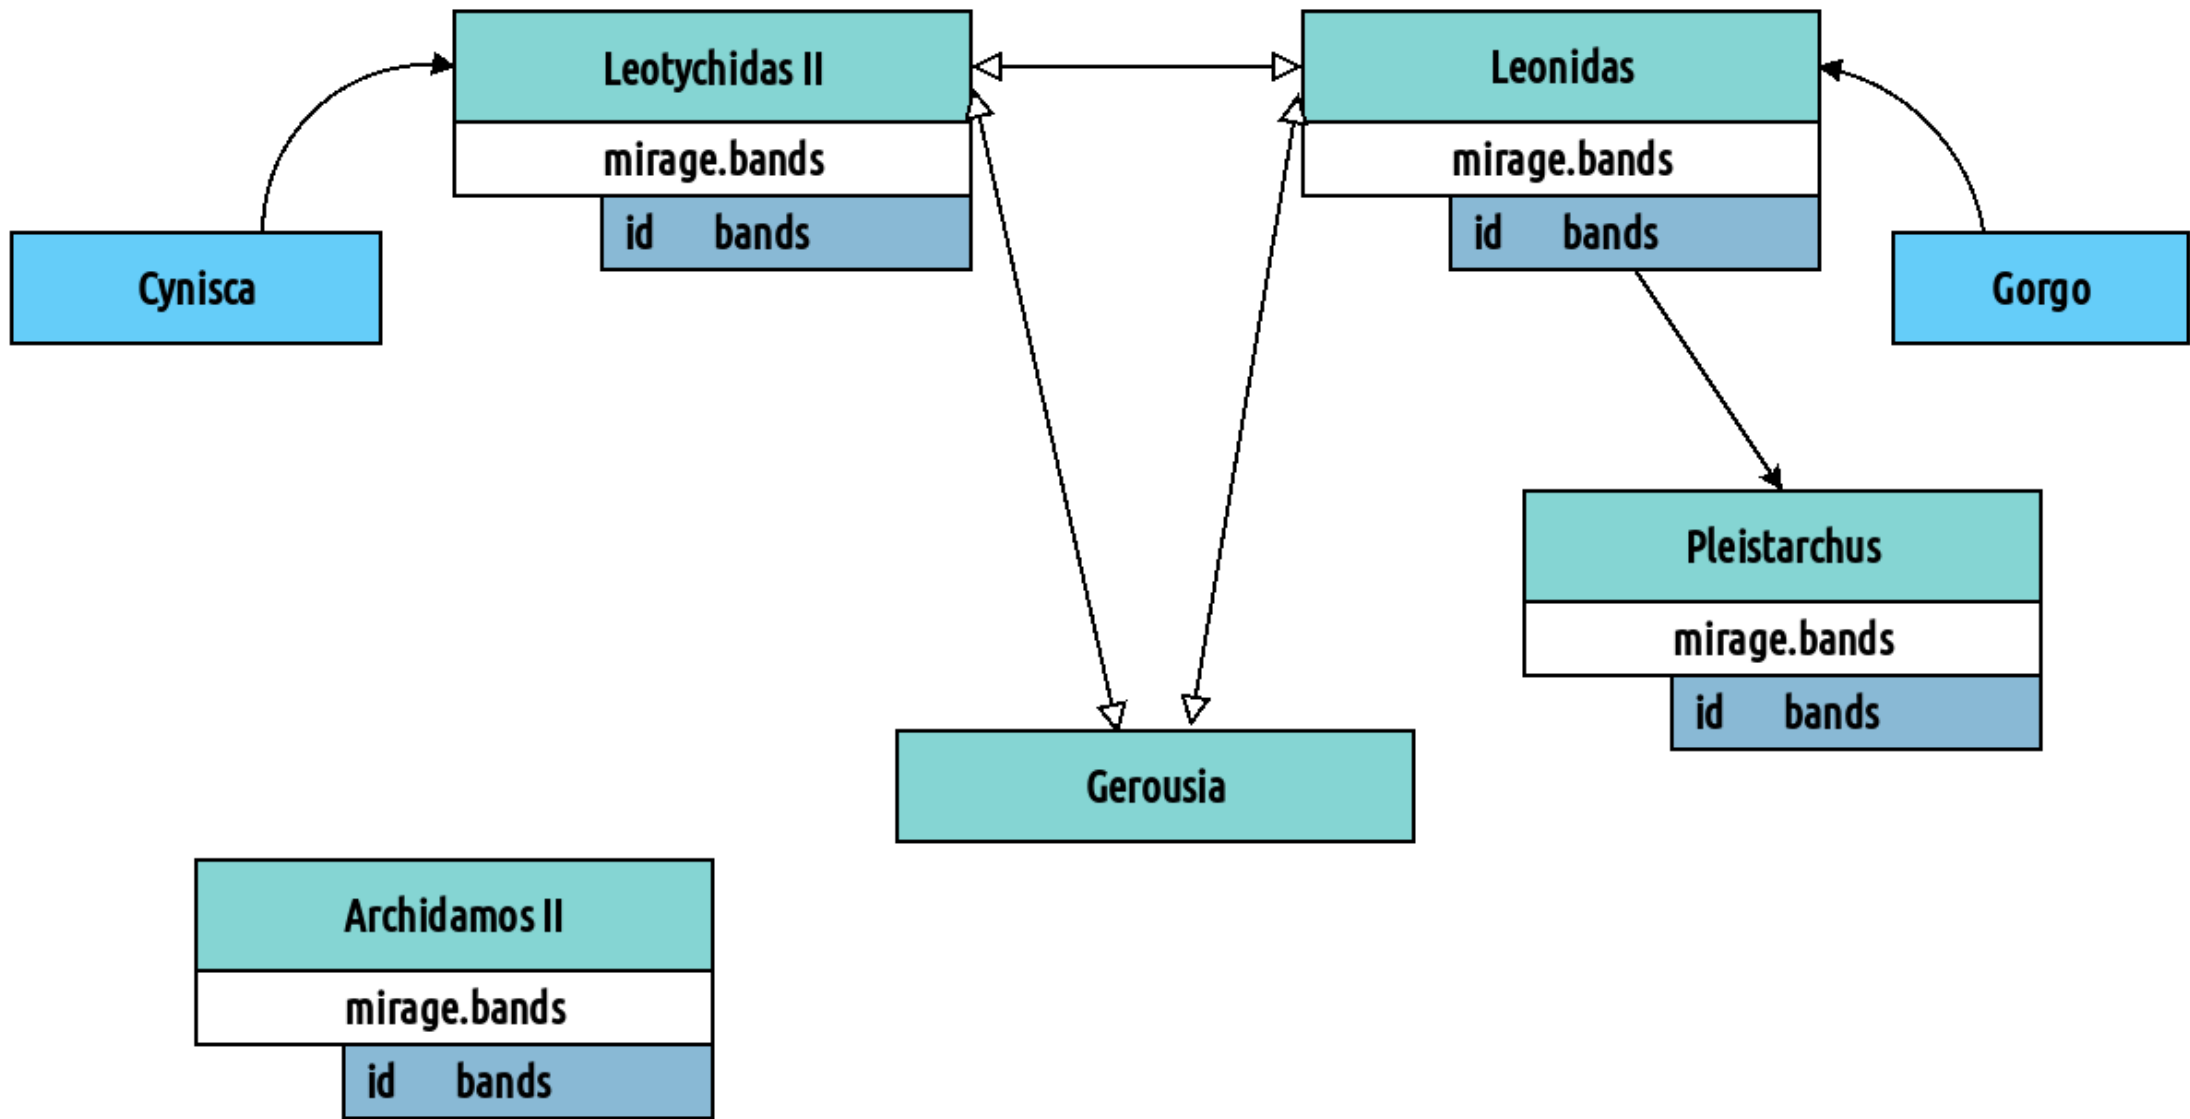


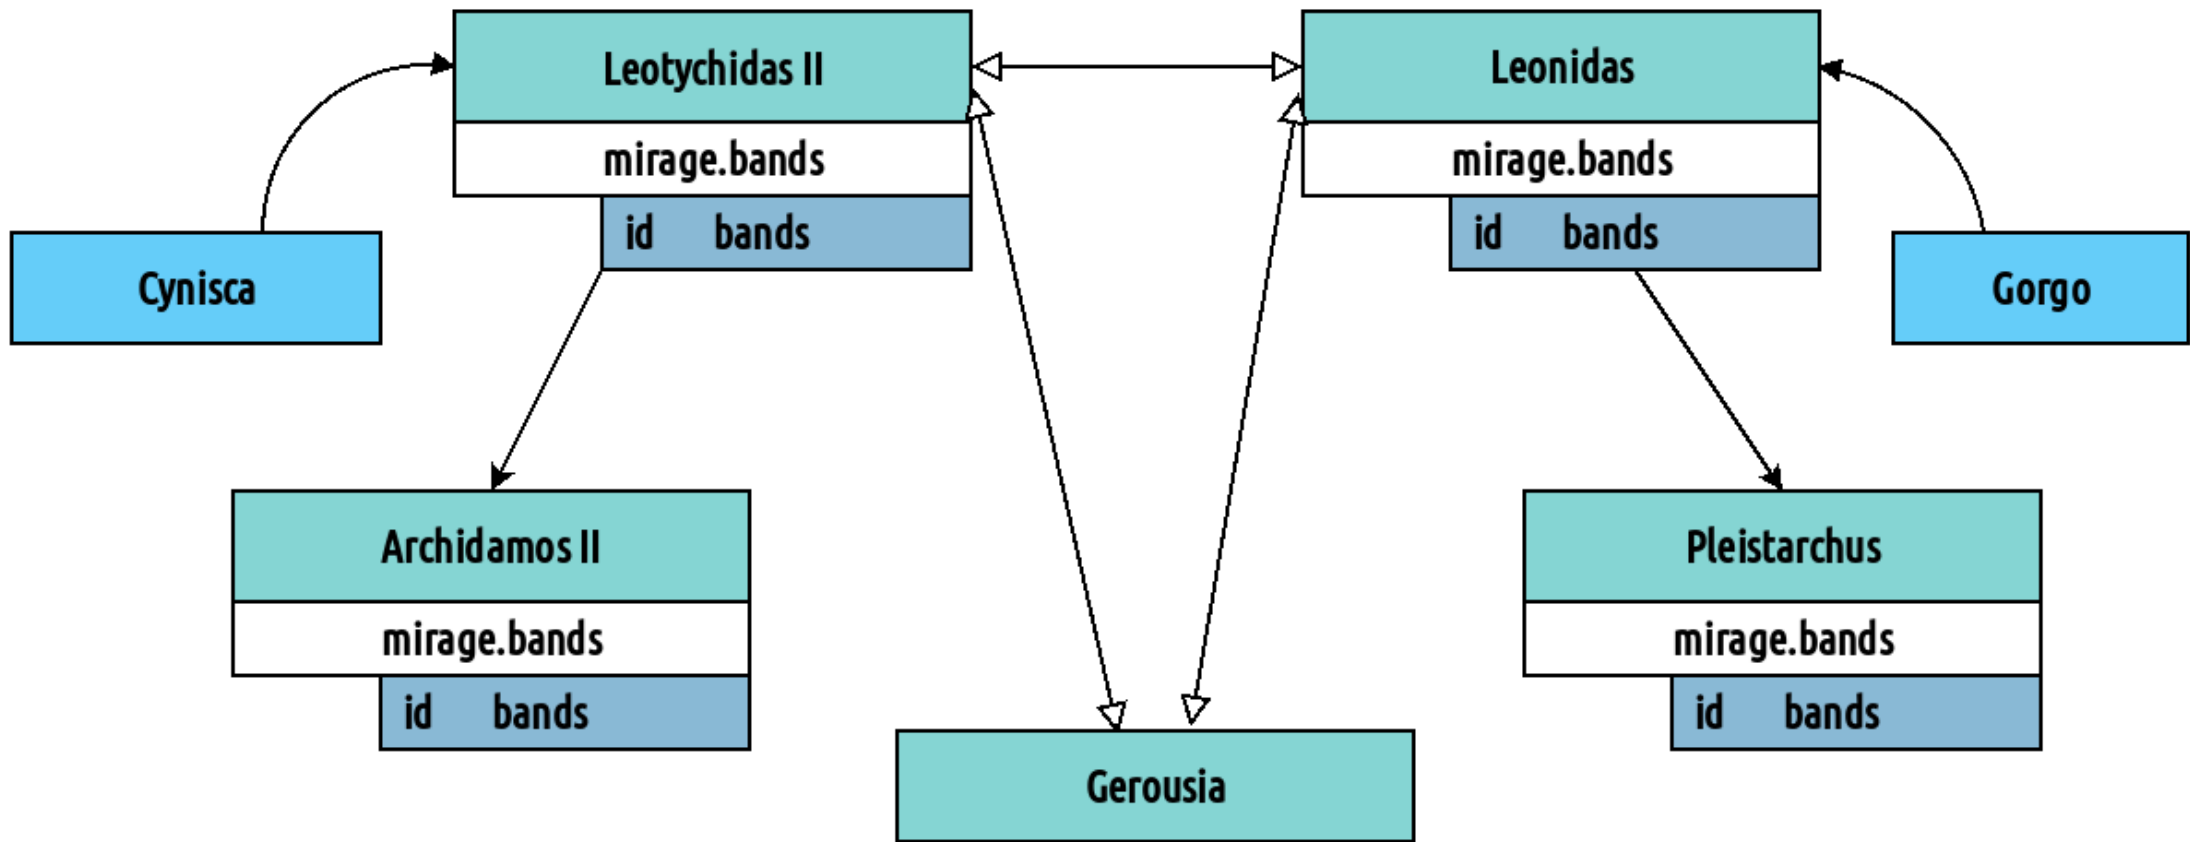


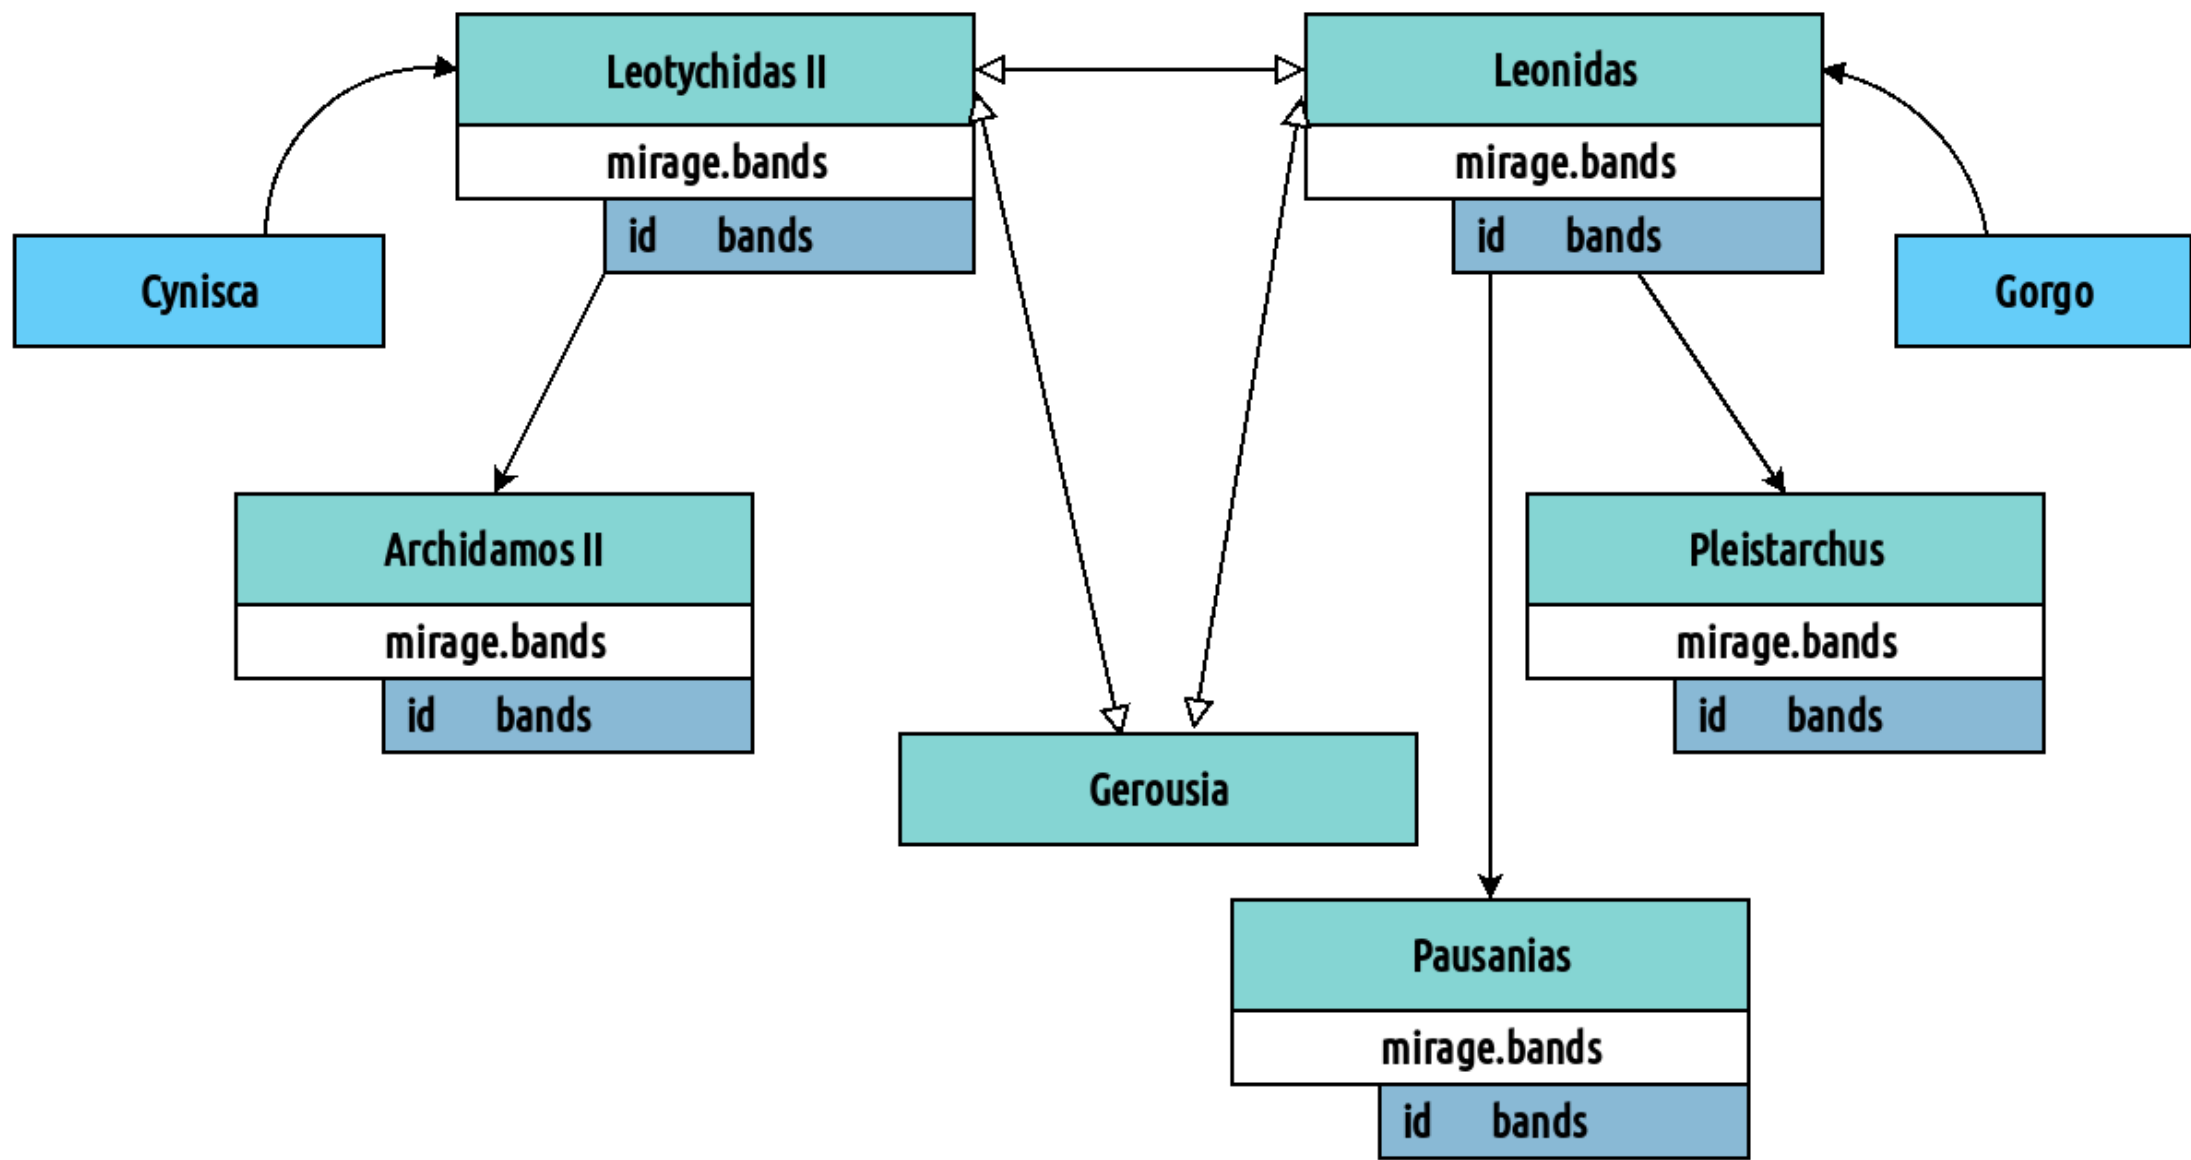


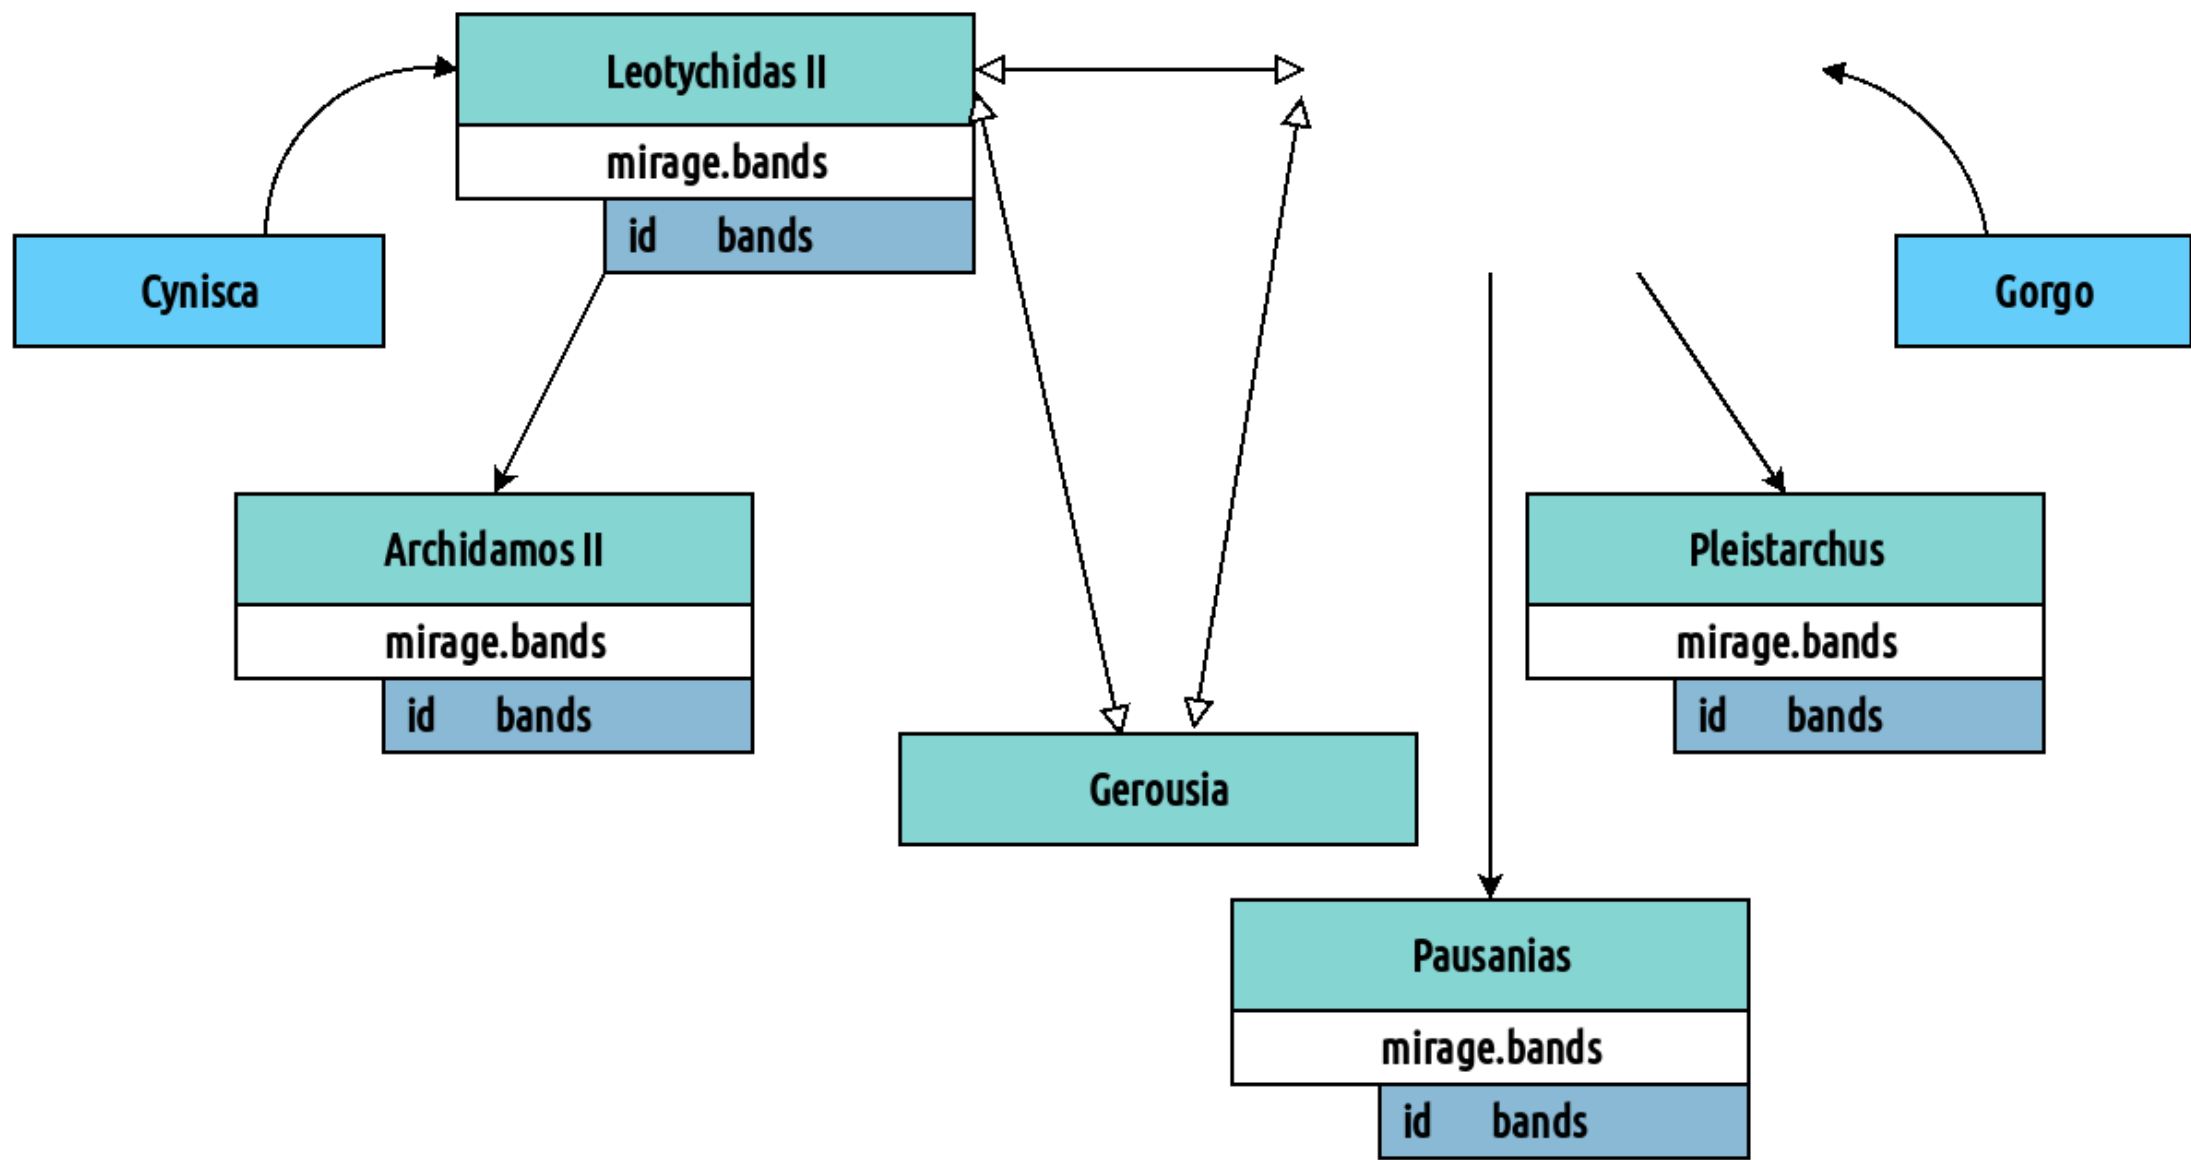


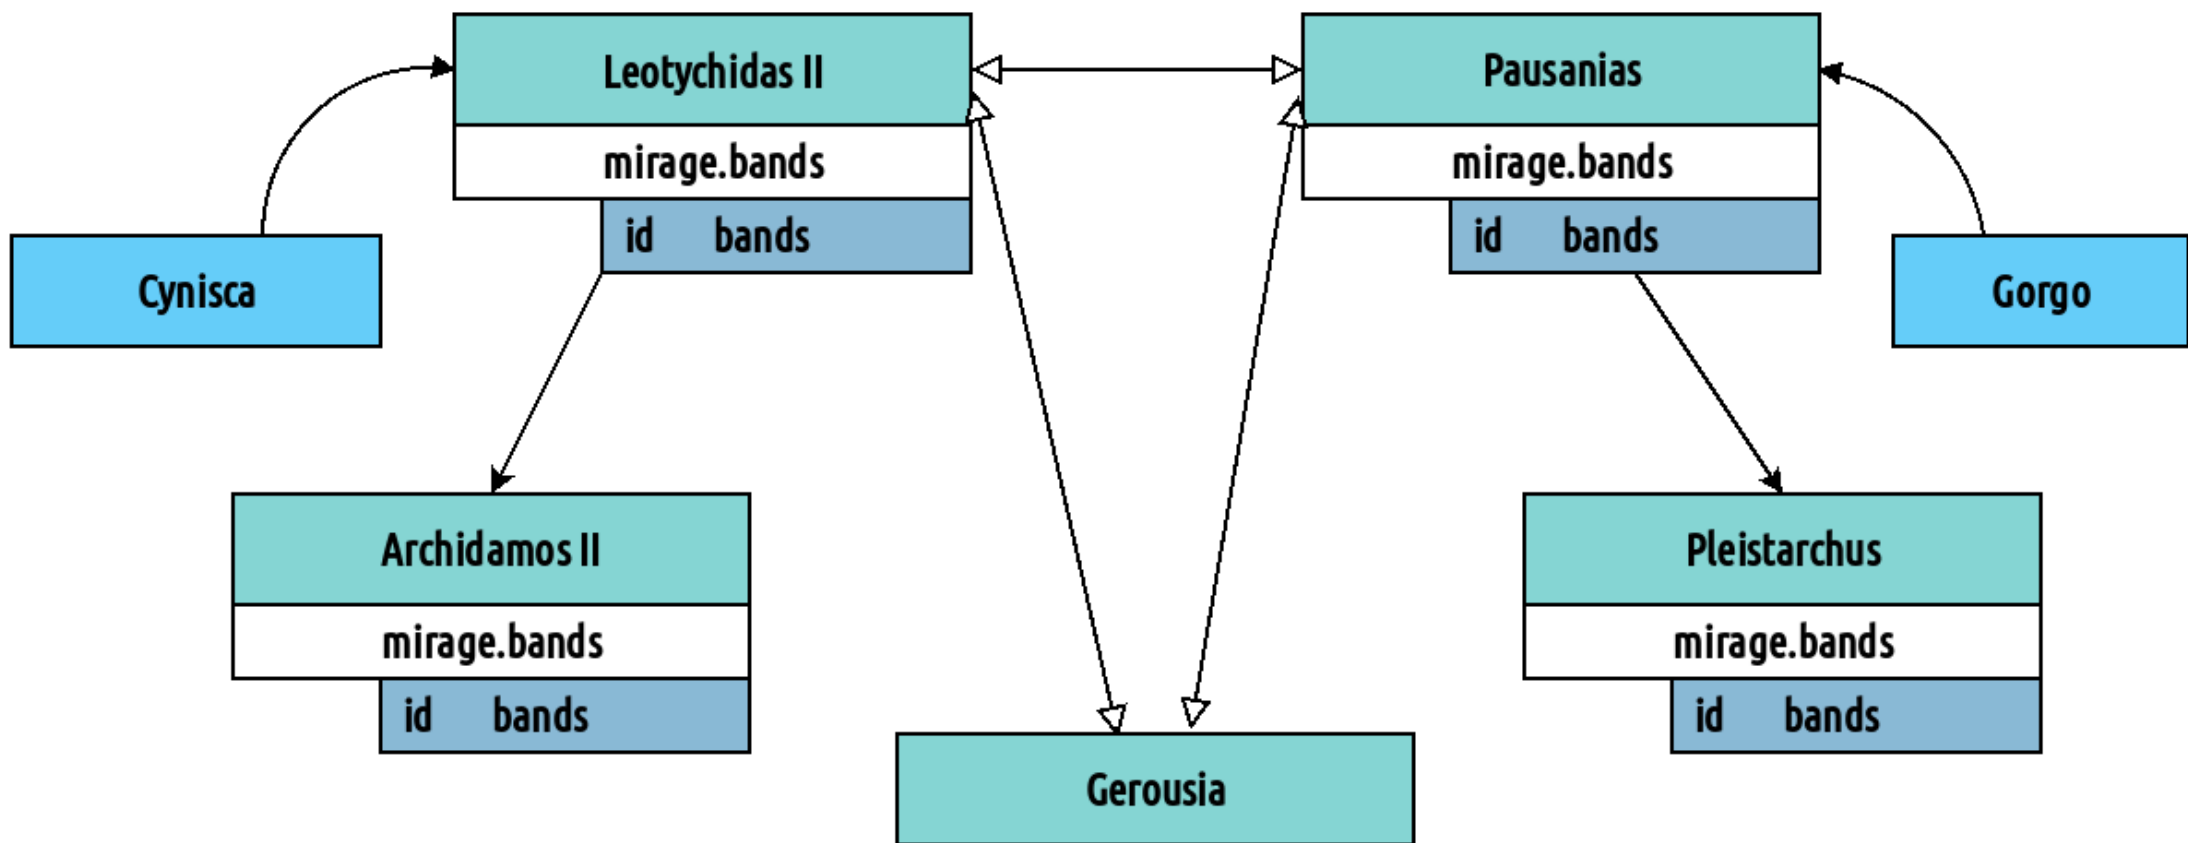


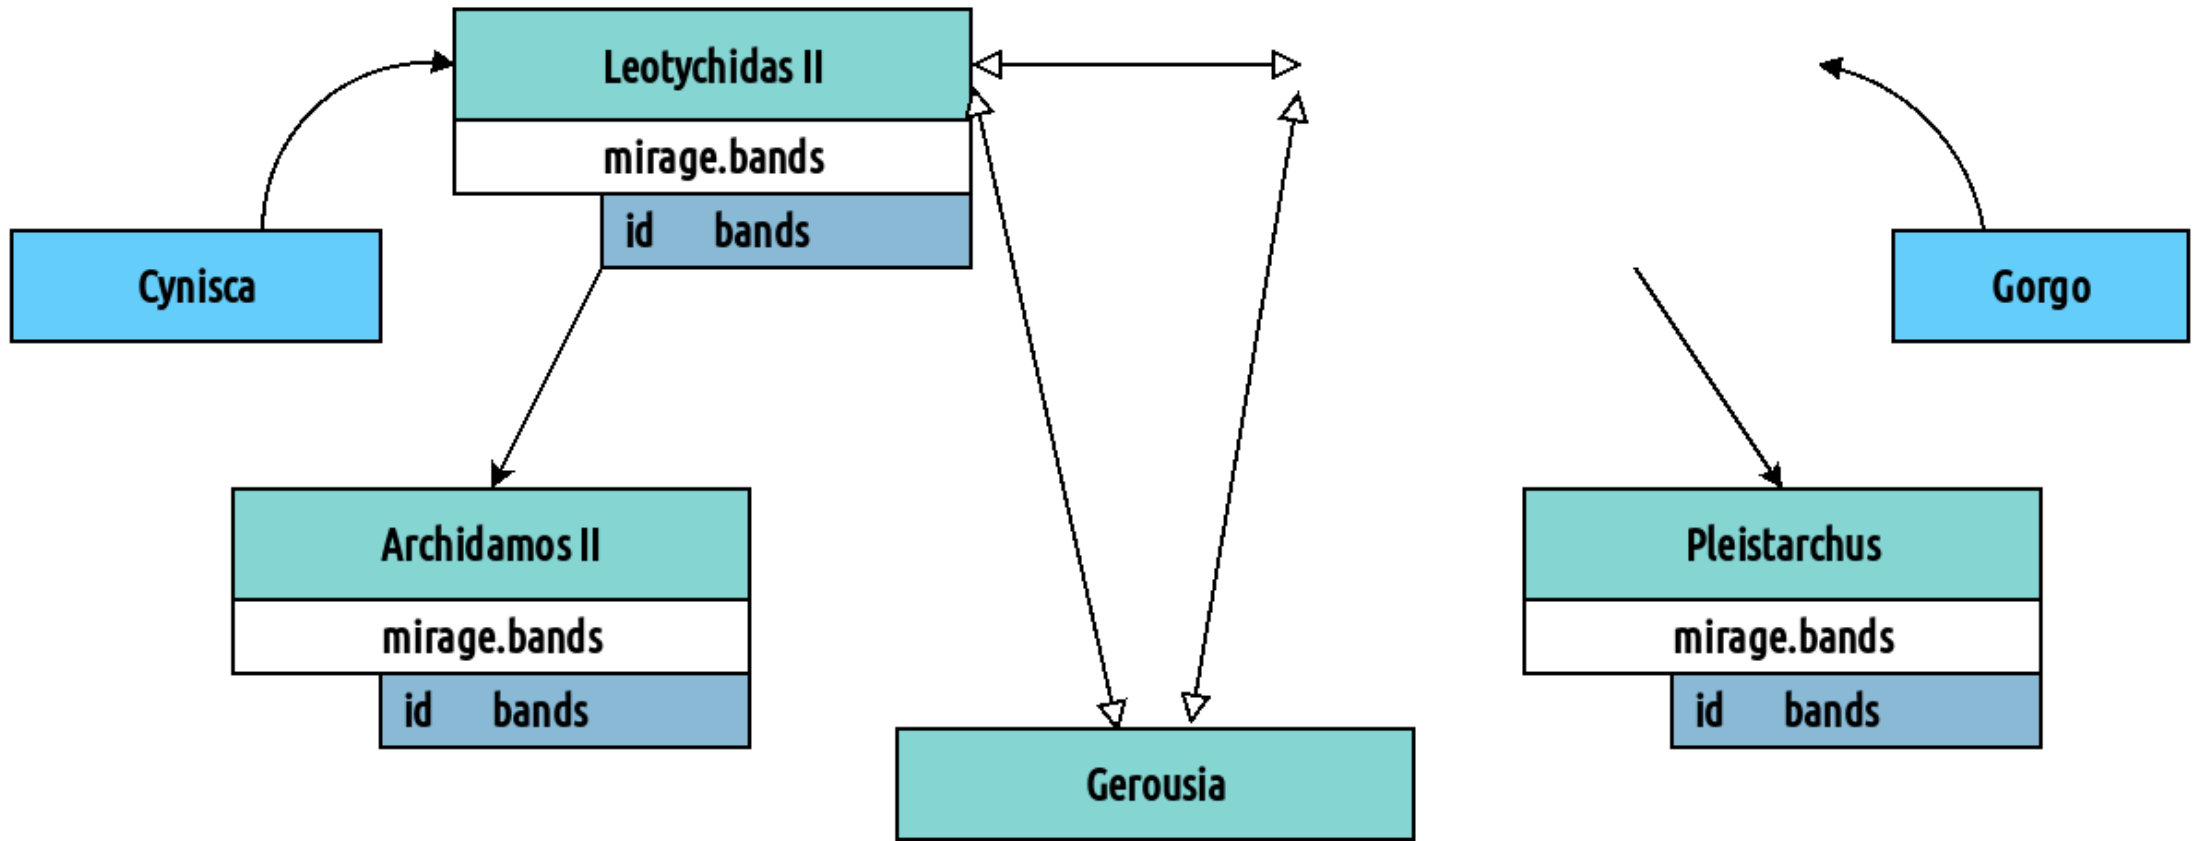


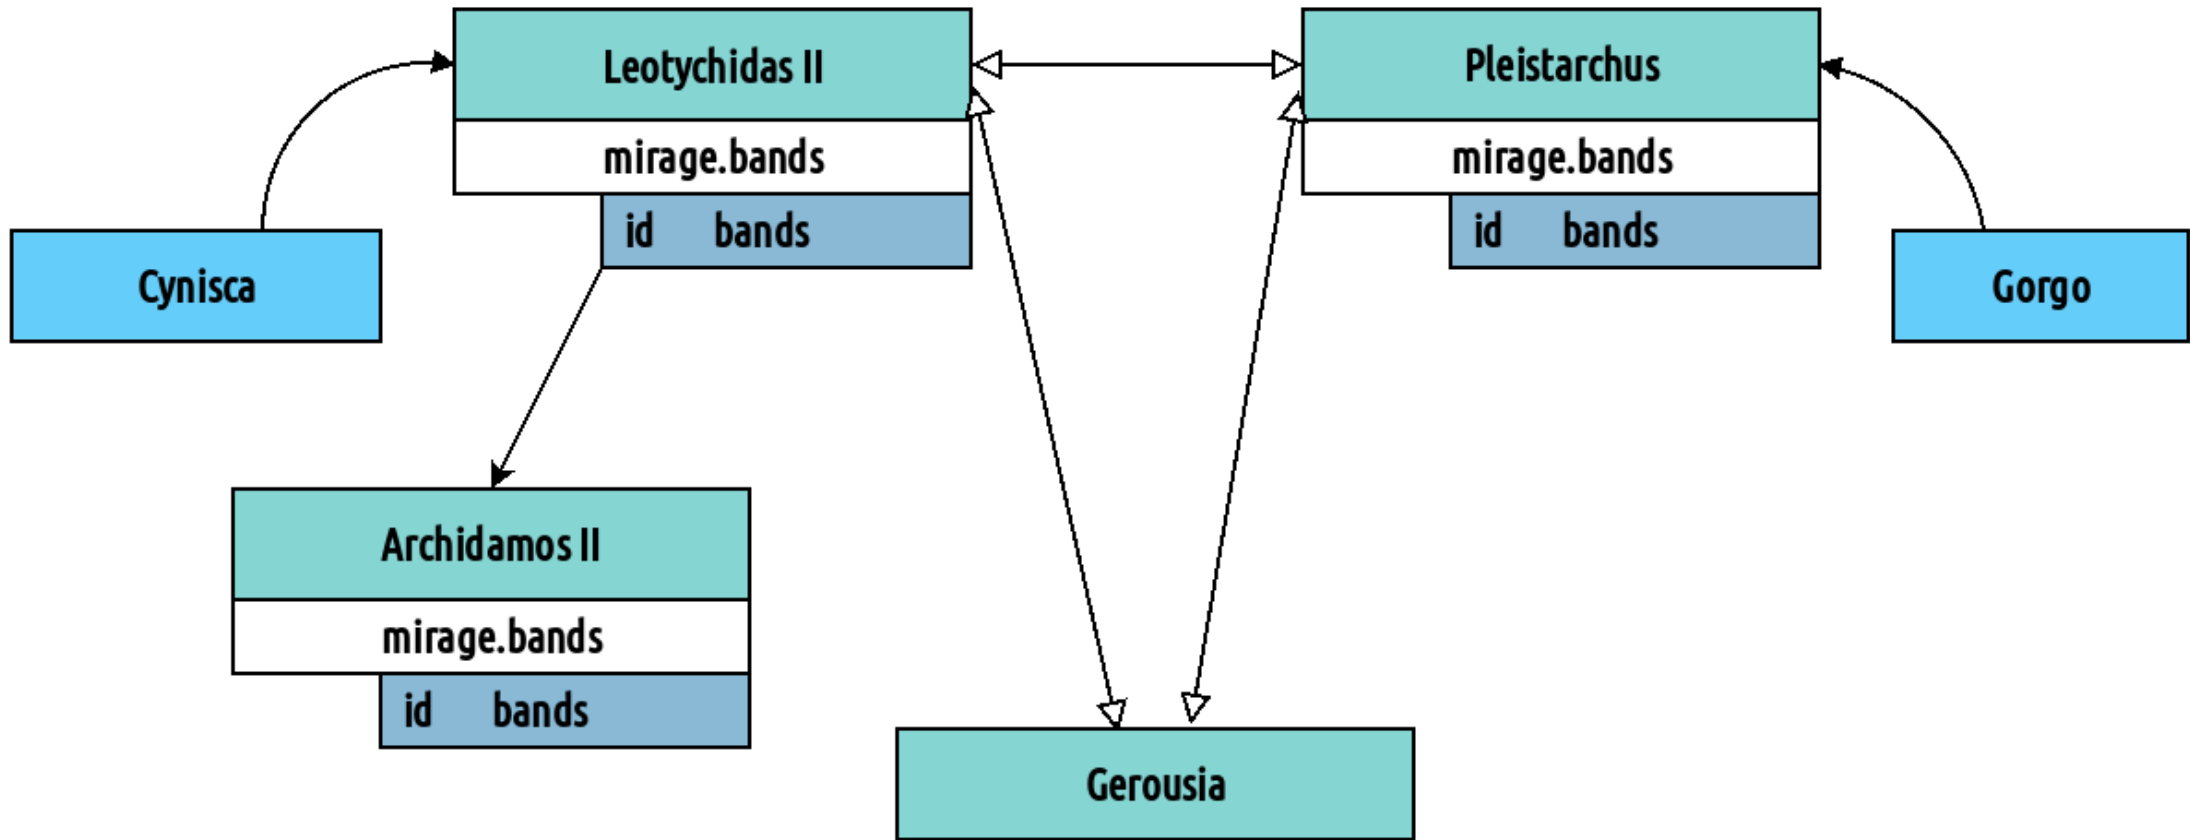


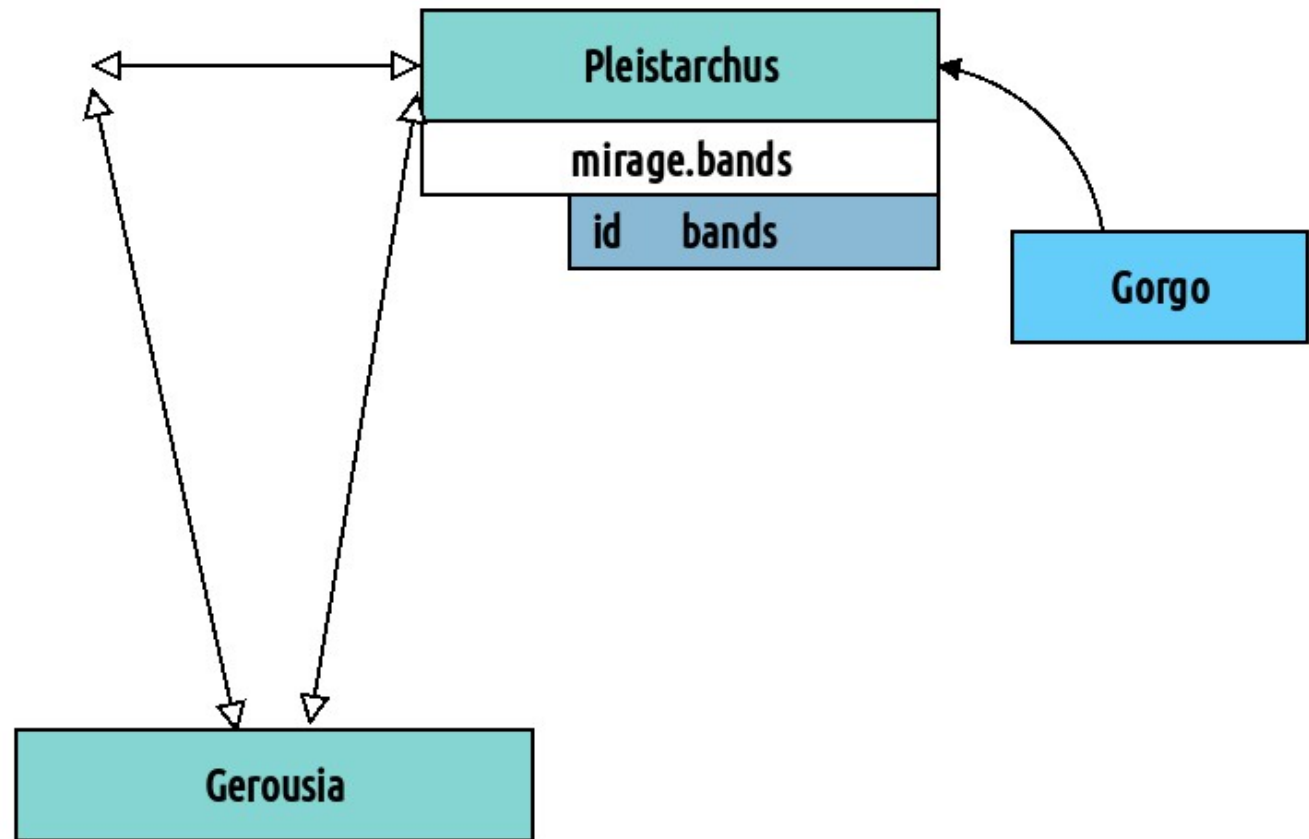
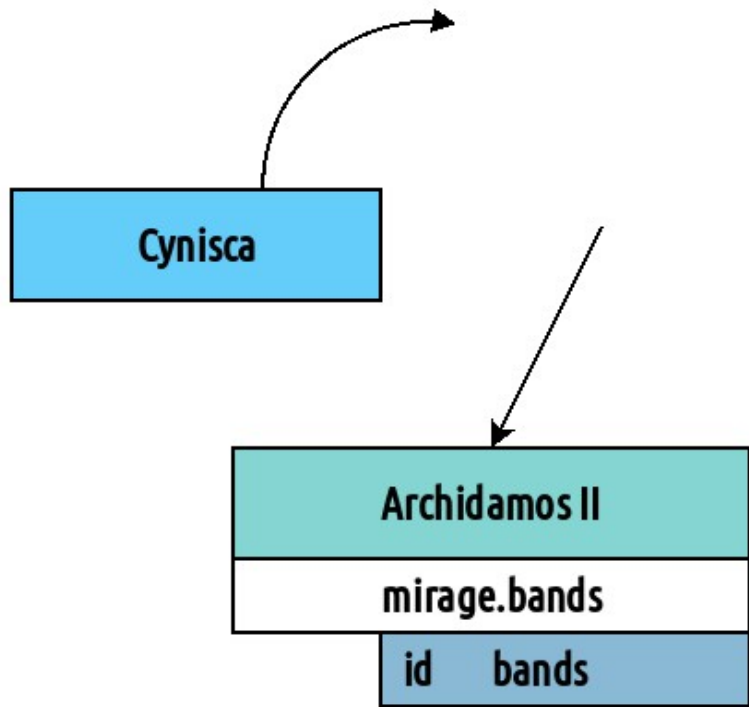


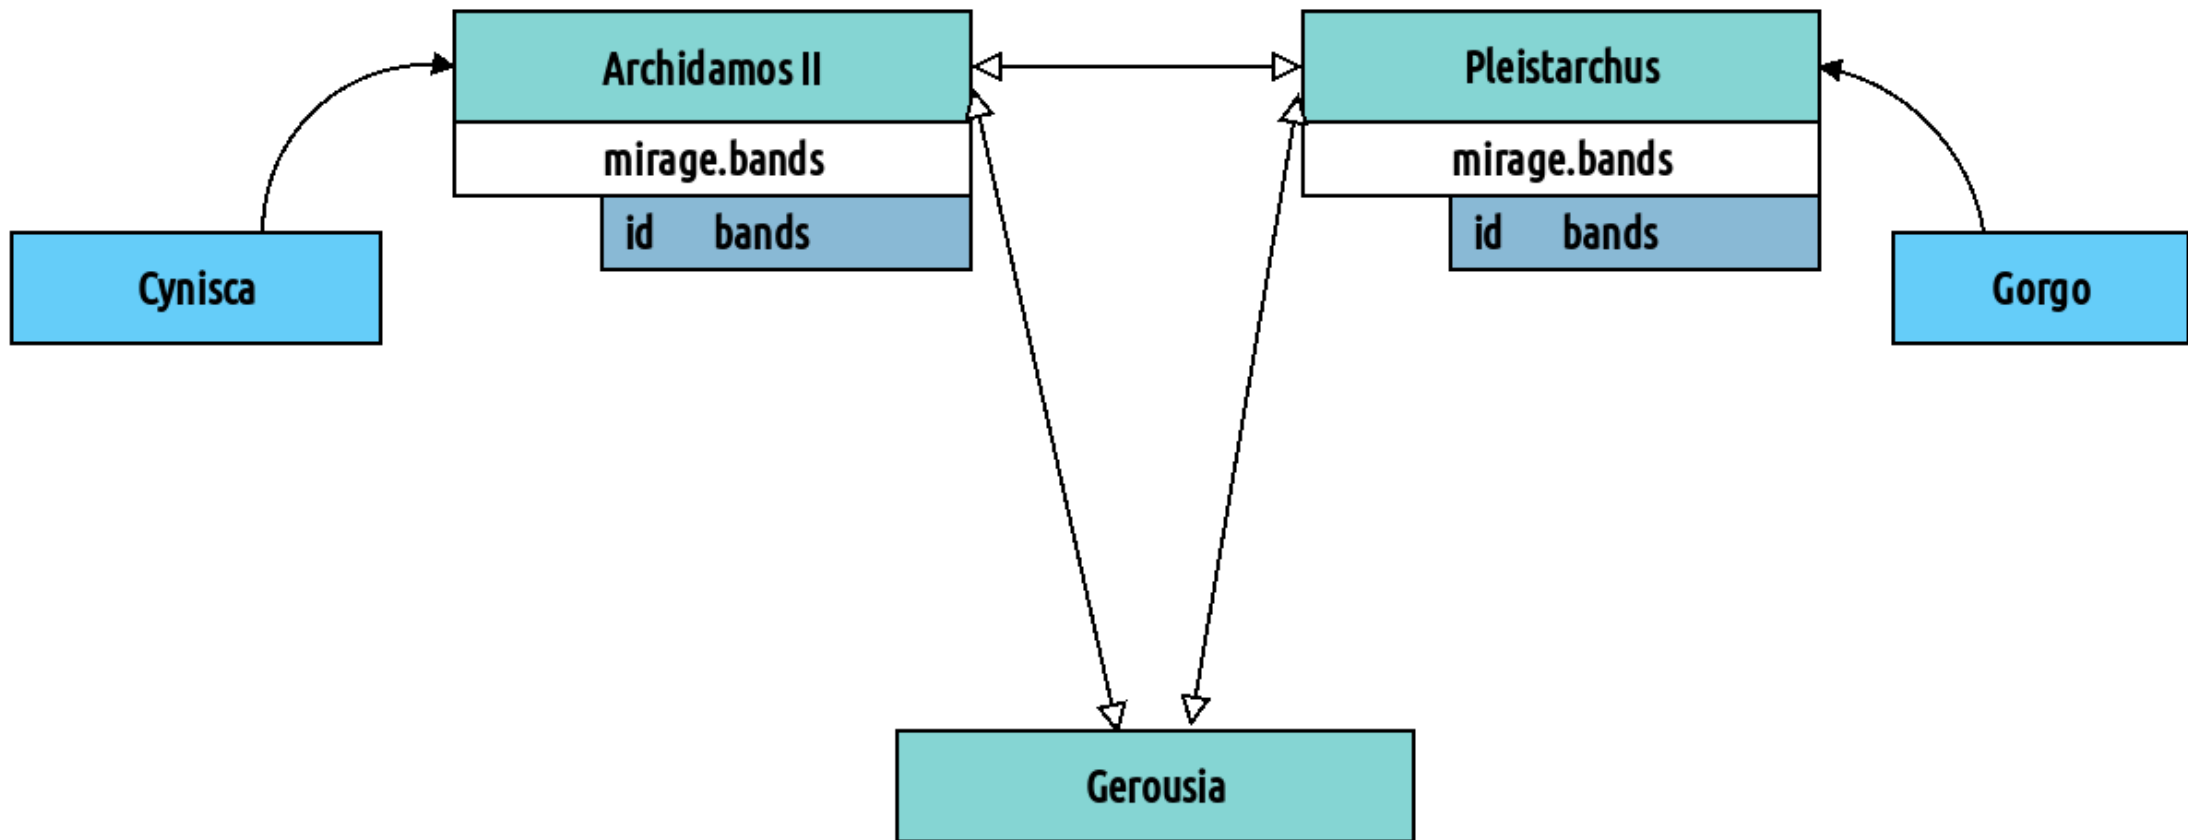












The Battle of Thermopylae

No chance of success

There was a disaster recovery plan



The Battle of Thermopylae

No chance of success

There was a disaster recovery plan

It was a moral victory that has last for many years



The Battle of Thermopylae

480 BCE Greco-Persian



The Battle of Thermopylae

480 BCE Greco-Persian

Gallic, Roman, Balkan, Byzantine, Bulgarian...



The Battle of Thermopylae

480 BCE Greco-Persian

Gallic, Roman, Balkan, Byzantine, Bulgarian...

1821 Greek Independence – 48 against 8000



The Battle of Thermopylae

480 BCE Greco-Persian

Gallic, Roman, Balkan, Byzantine, Bulgarian...

1821 Greek Independence – 48 against 8000

1941 WWII – Sabotage of the viaduct



Disaster Recovery Plan
is not only about the servers



High Availability
is not only about the servers



Wrapping up



Scale-out

High Availability

Geo Redundancy



Atomicity
Consistency
Isolation
Durability

Consistency
Availability
Partition Tolerance



Spartan's Dual Kingship

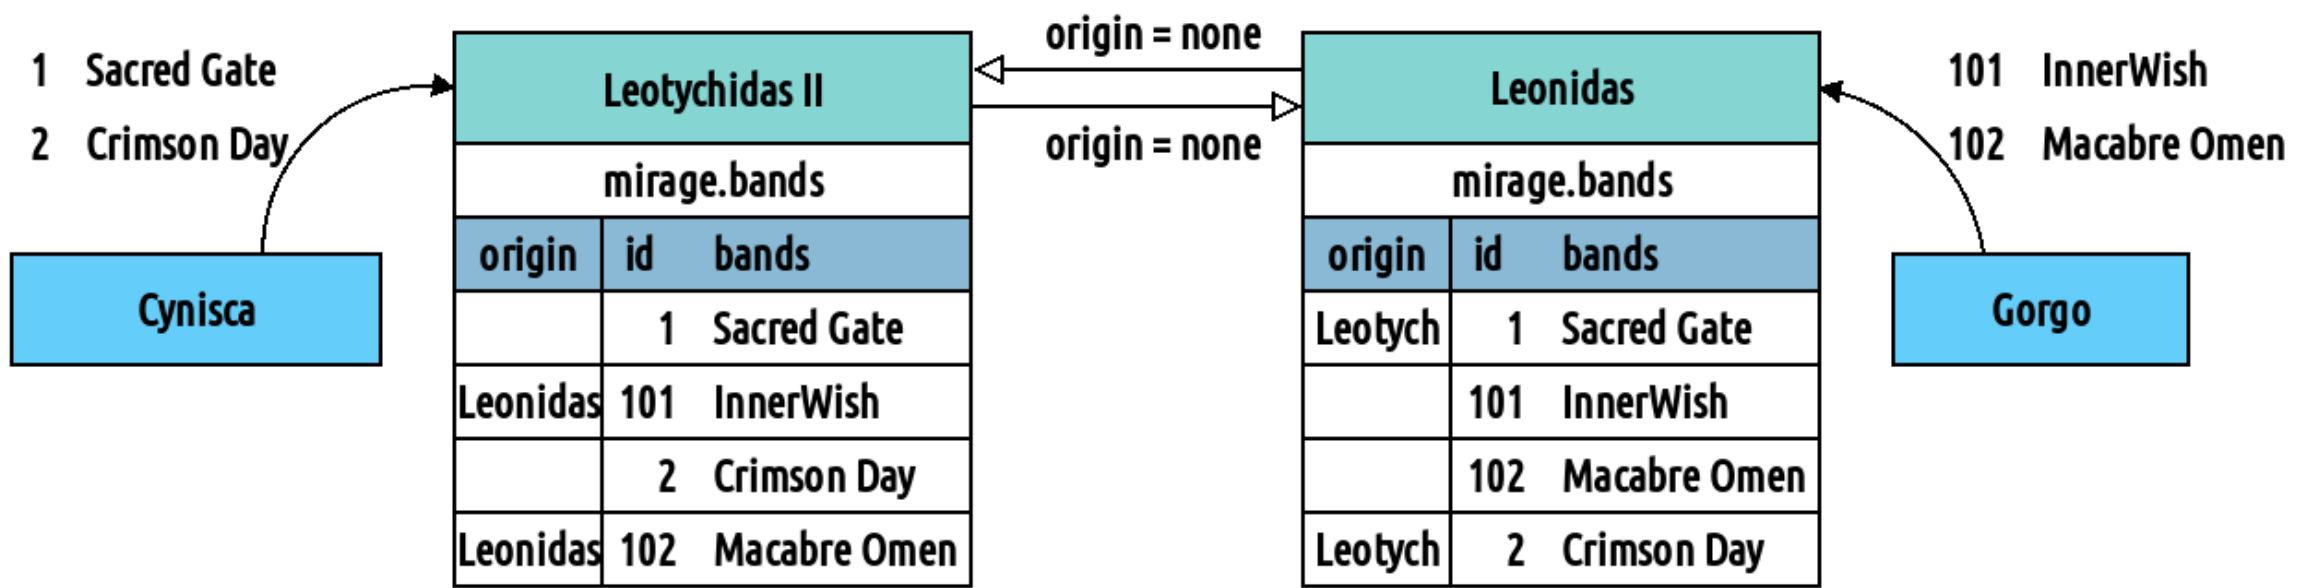



Euryontid

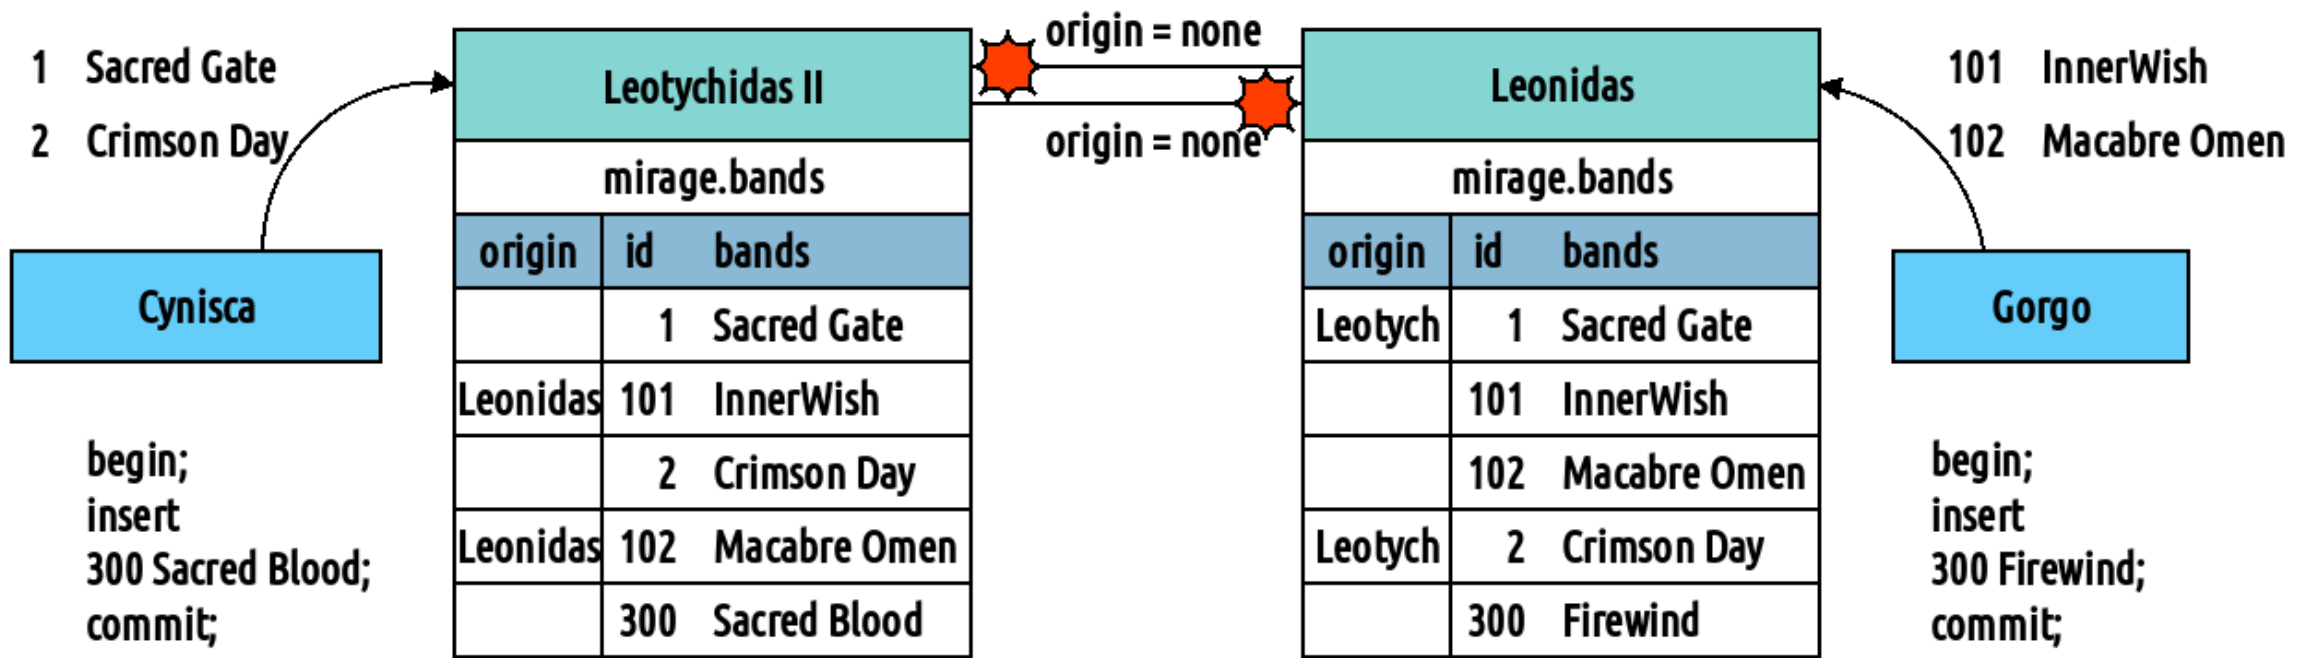


Agiad










Consistency

Every read receives the most recent write
or an error






	300	Sacred Blood	ts1
--	-----	--------------	-----

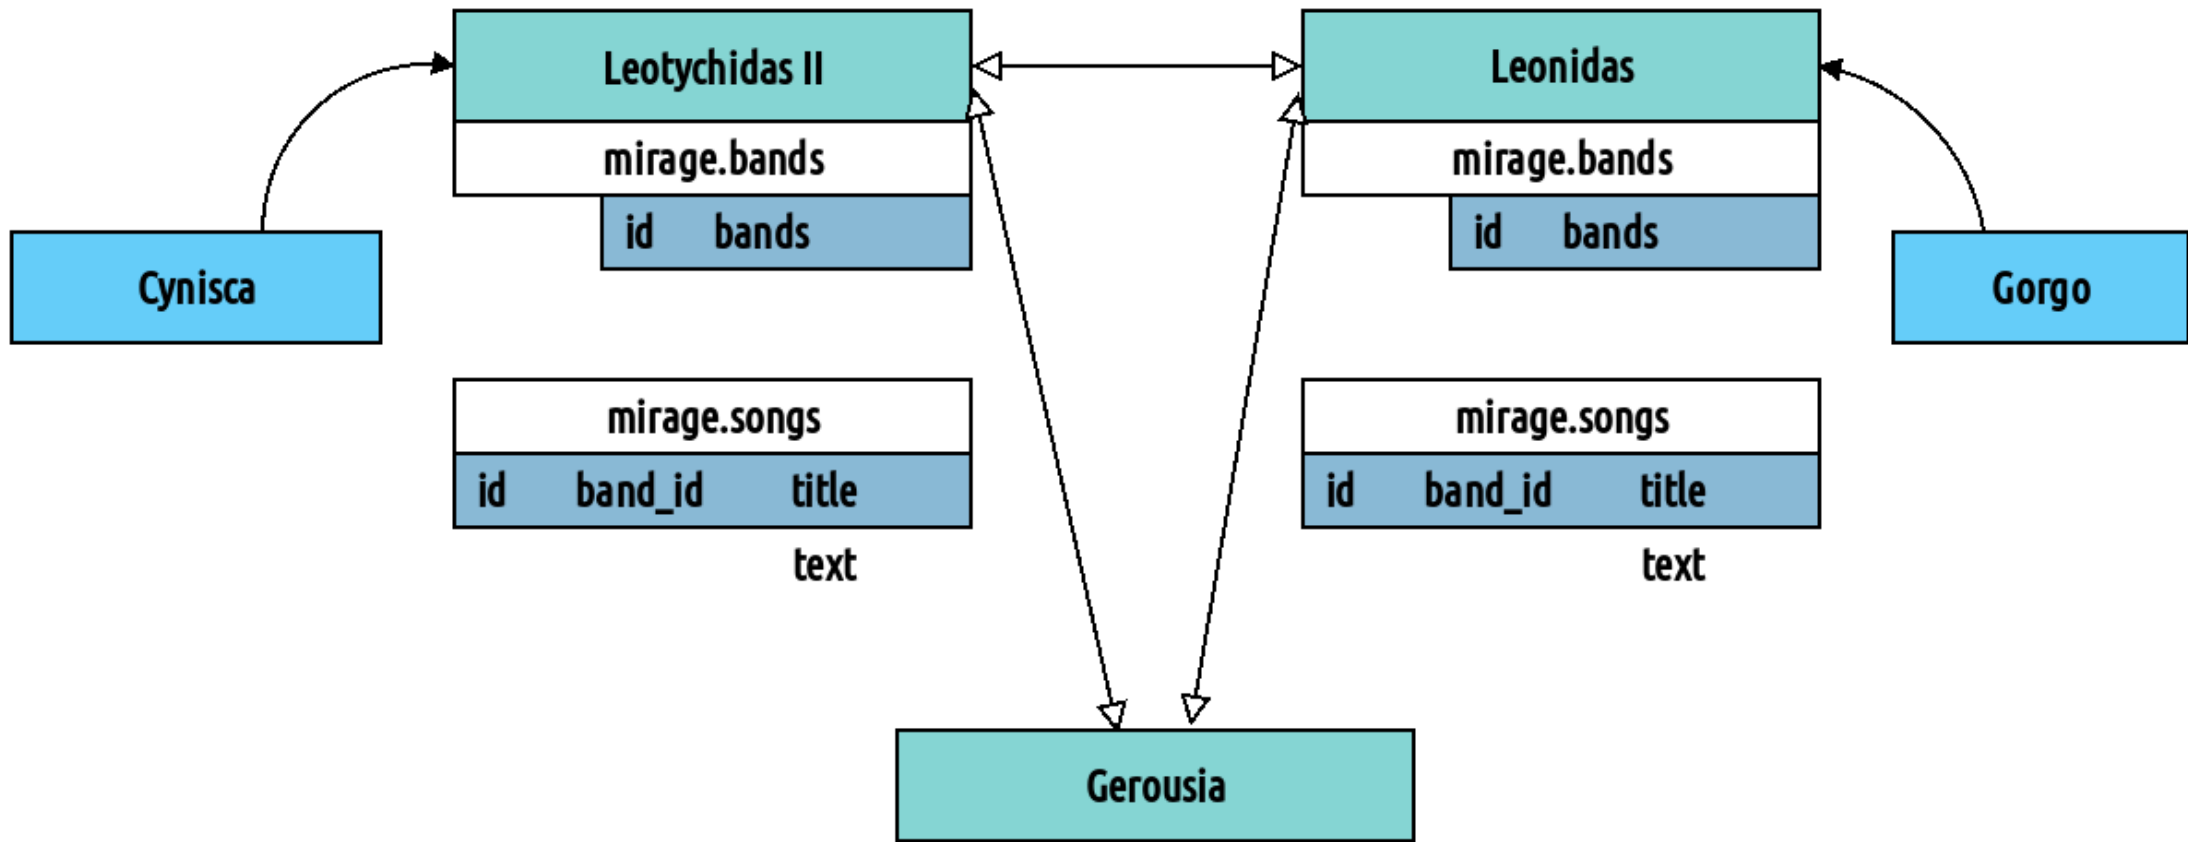
Leotych	300	Sacred Blood	ts1
---------	-----	--------------	-----


Spartan's Dual Kingship and the Gerousia



Having consensus is very useful

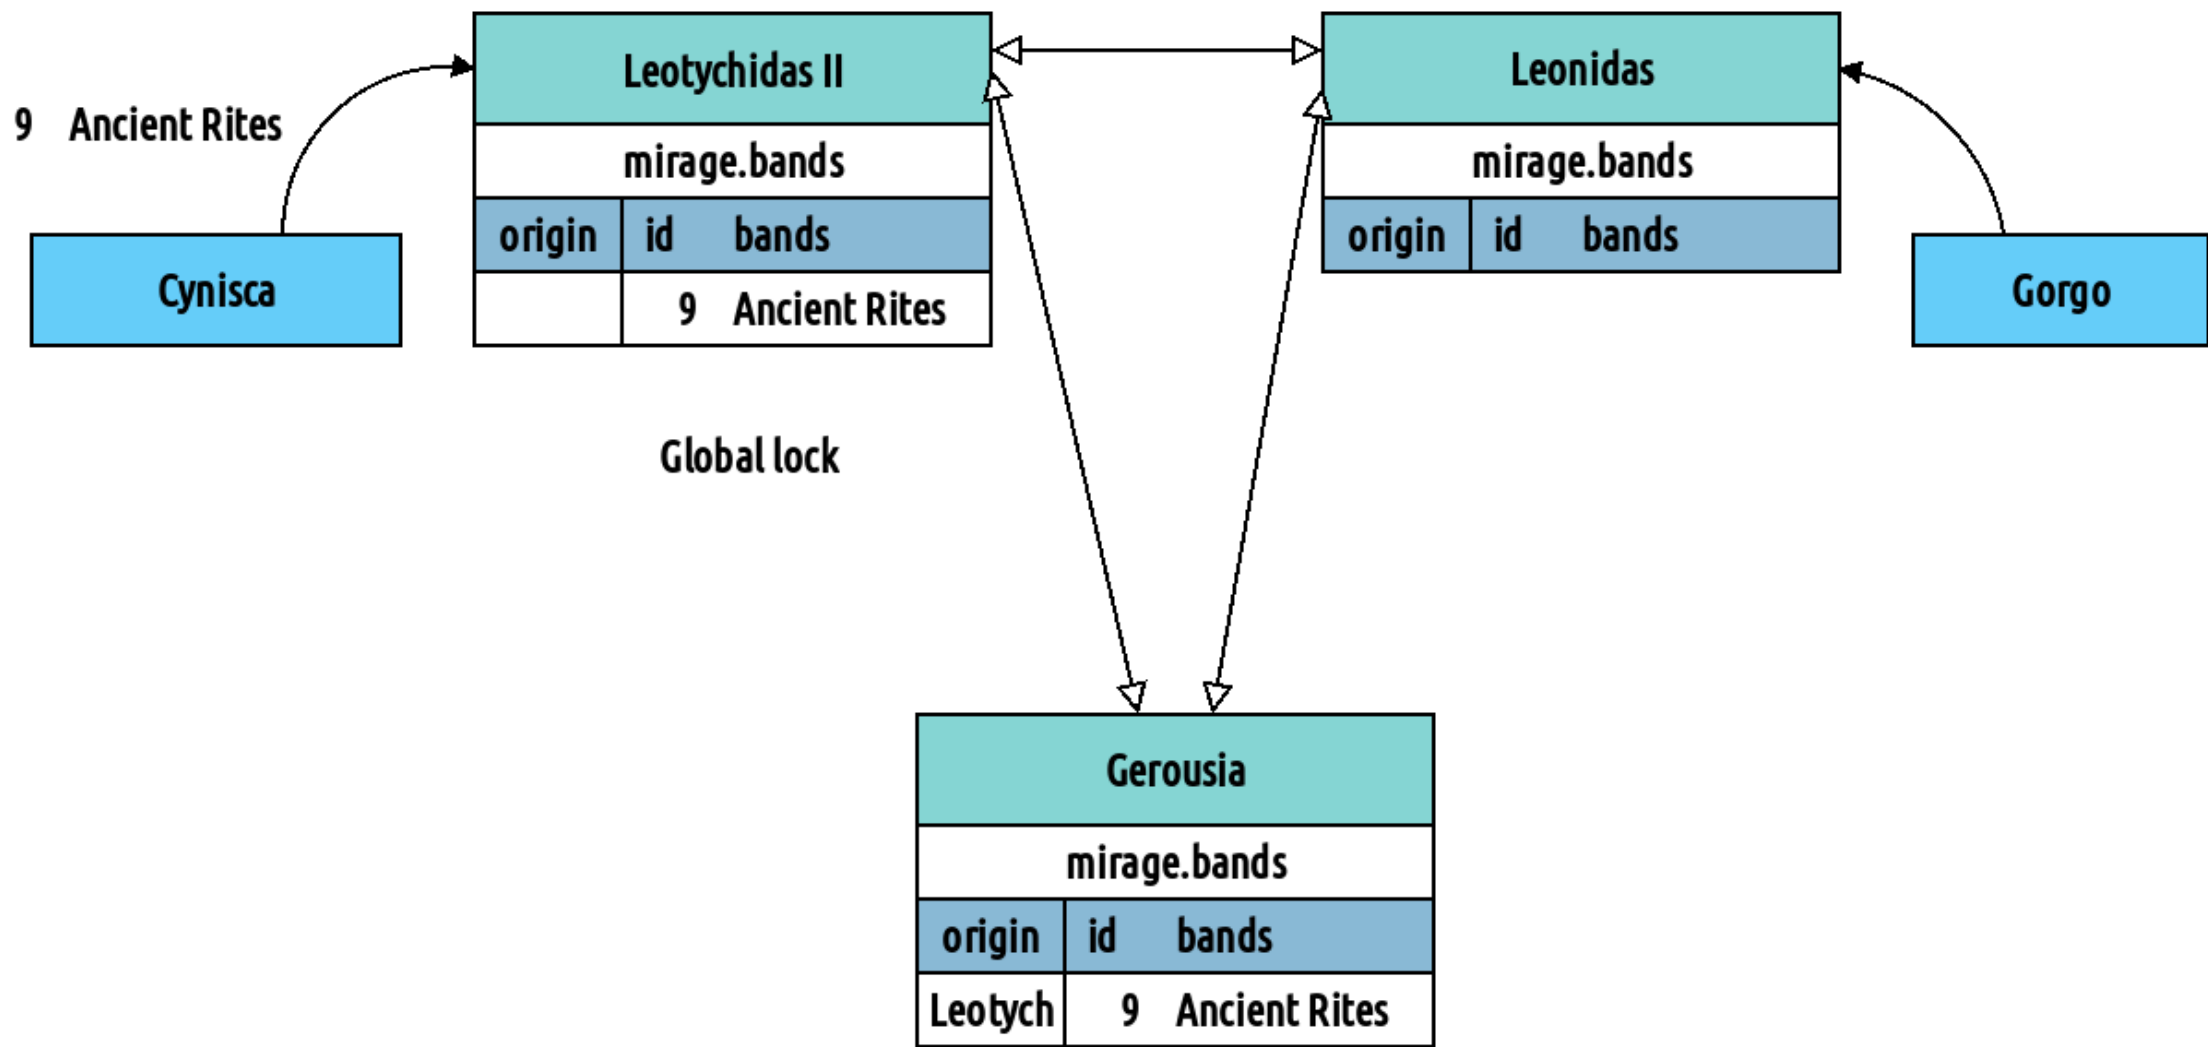






Reduce remote locking





9 Ancient Rites

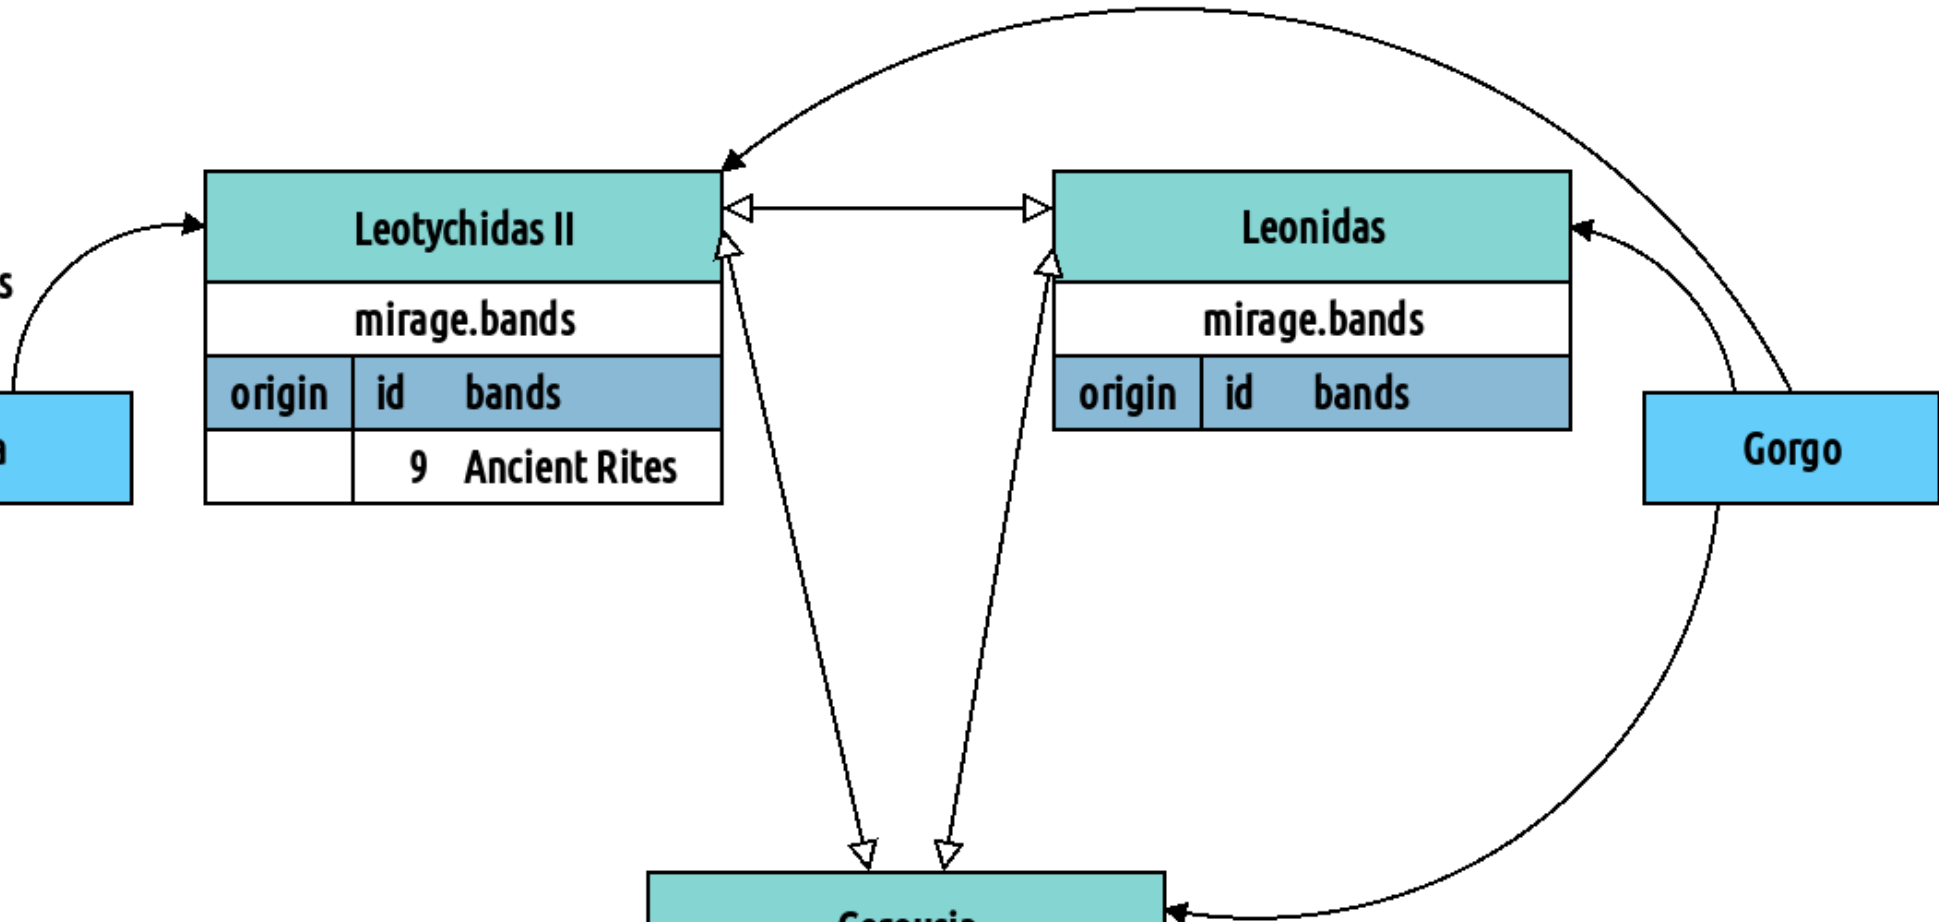
Cynisca

Leotychidas II		
mirage.bands		
origin	id	bands
	9	Ancient Rites

Leonidas		
mirage.bands		
origin	id	bands

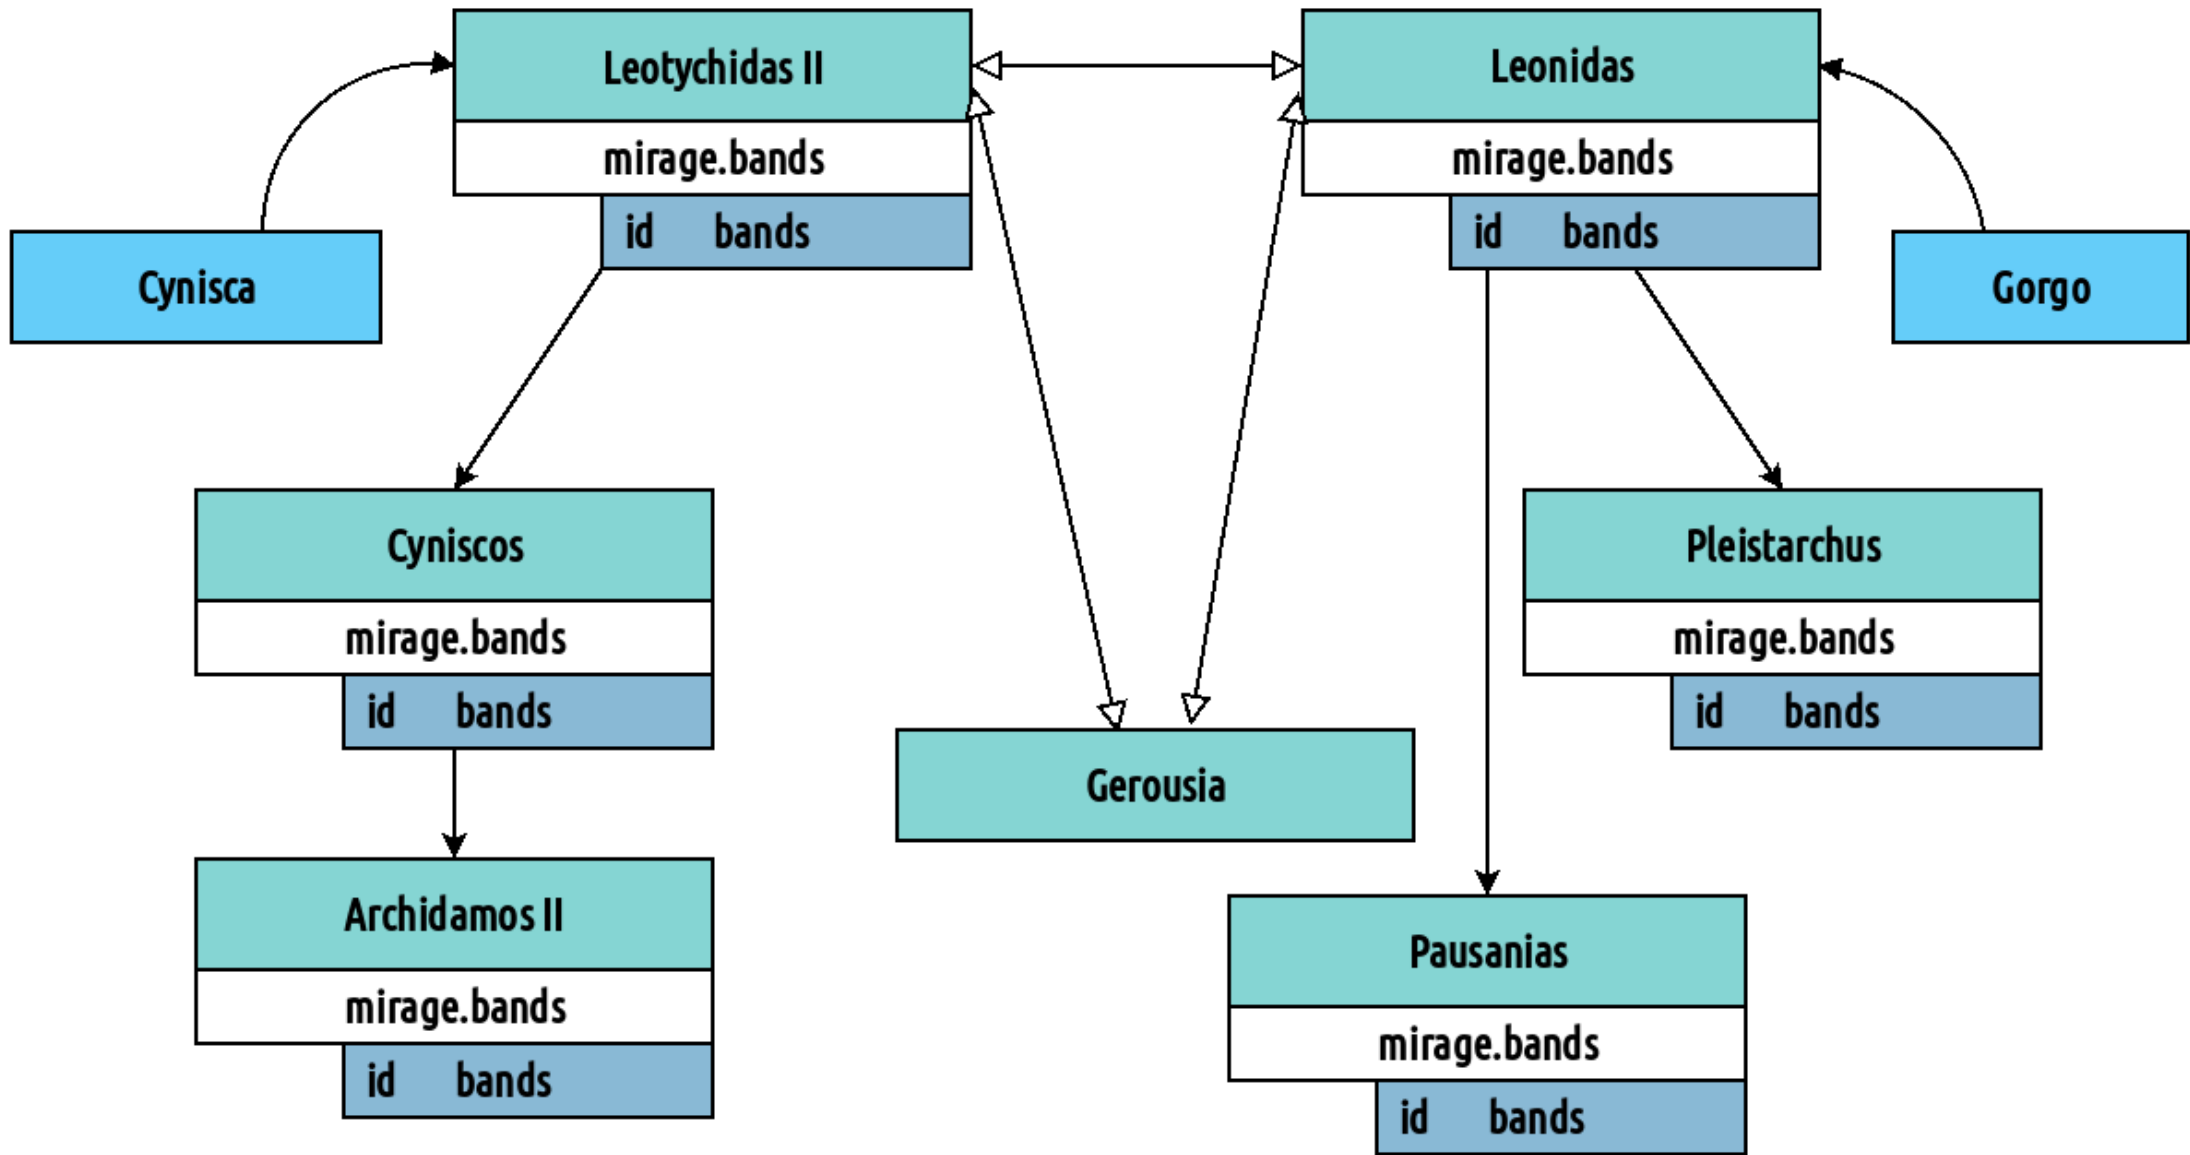
Gorgo

Gerousia		
mirage.bands		
origin	id	bands
Leotych	9	Ancient Rites



Disaster Recovery Mixing Logical and Physical Replication





Beware of the Mirage of Active-Active



Test your assumptions



Design with Active-Active in mind



Thank you

