

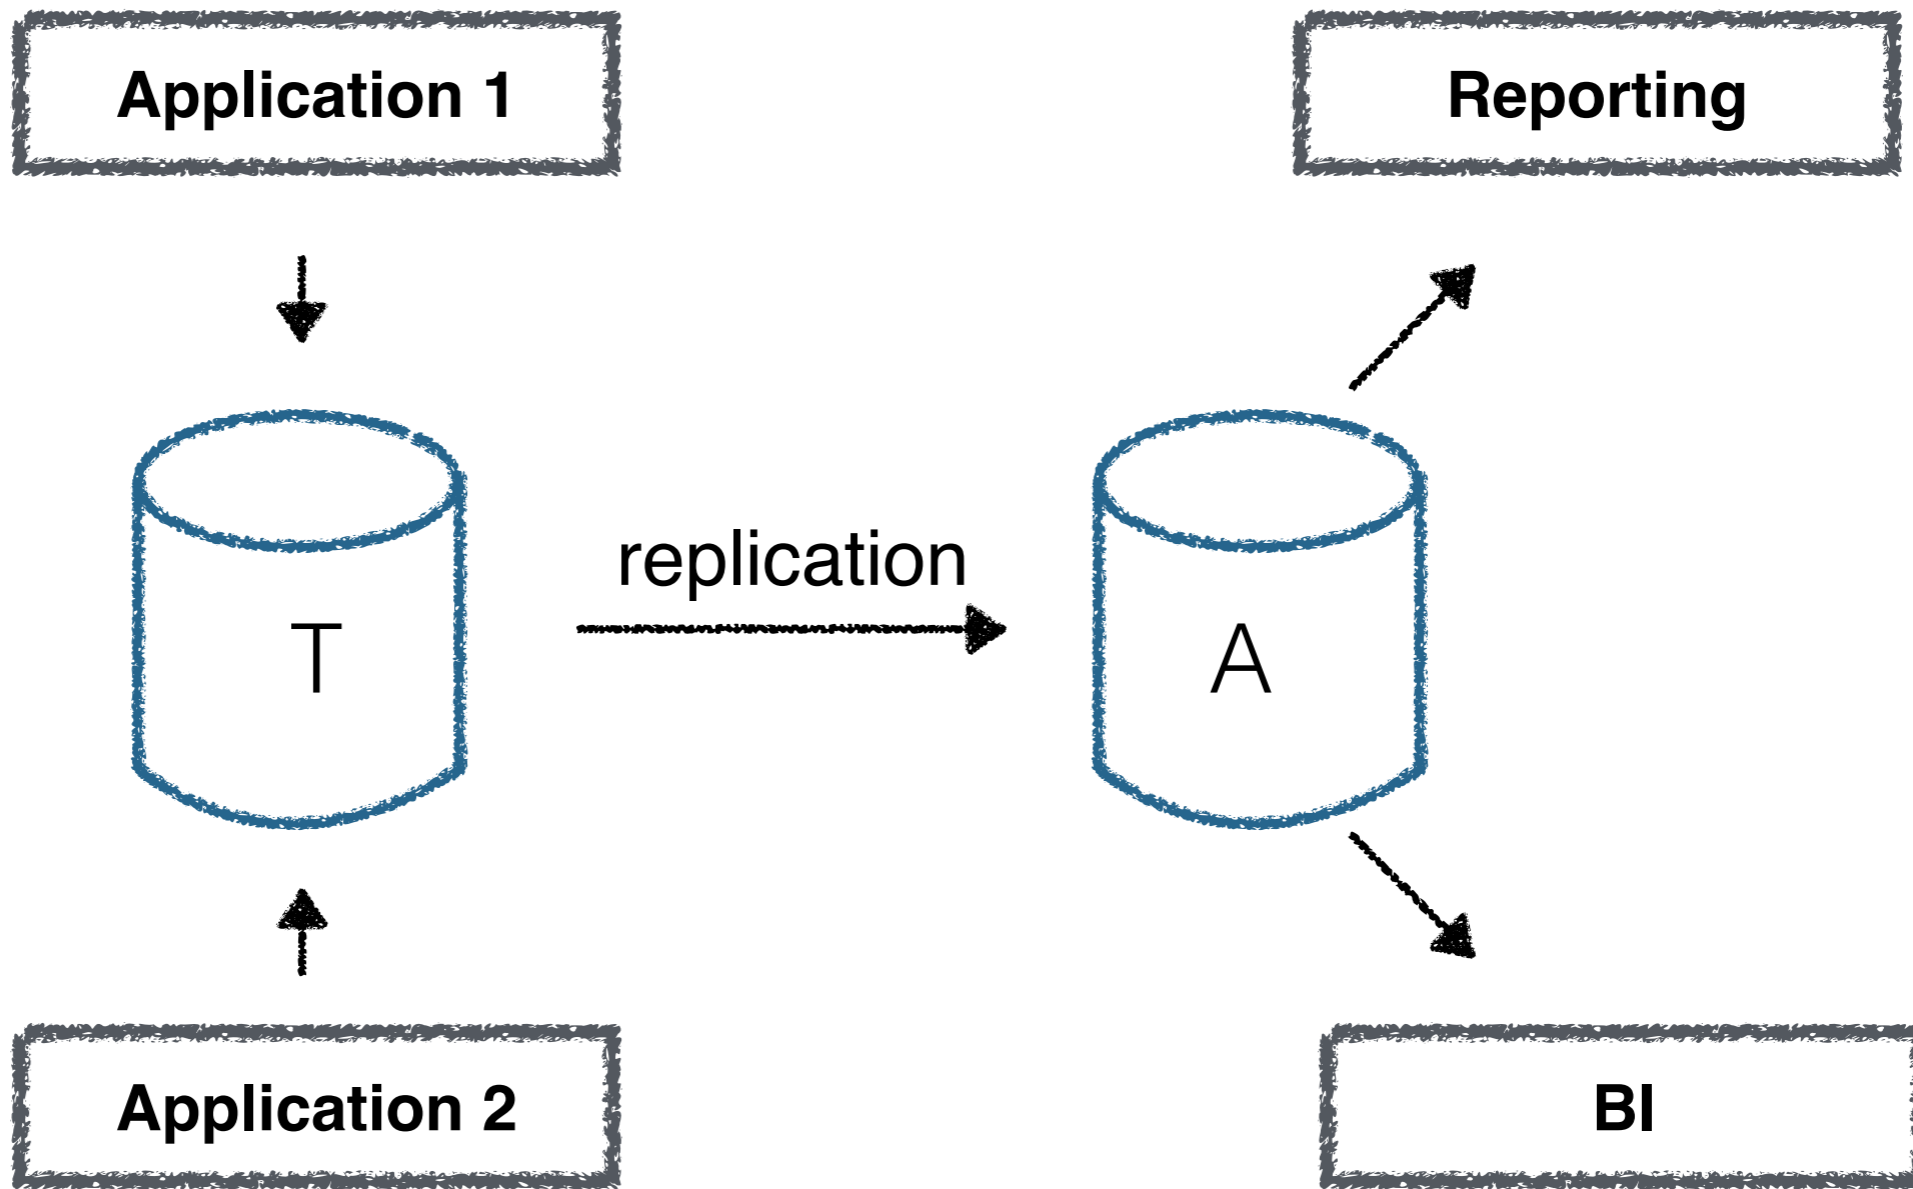
Surviving sharding

Aliaksandr Aliashkevich
Prague, 2015

Database Optimisation

- Query Optimisation
- Disk I/O
- CPU
- Memory

Scaling Reads




MPP Style



Application



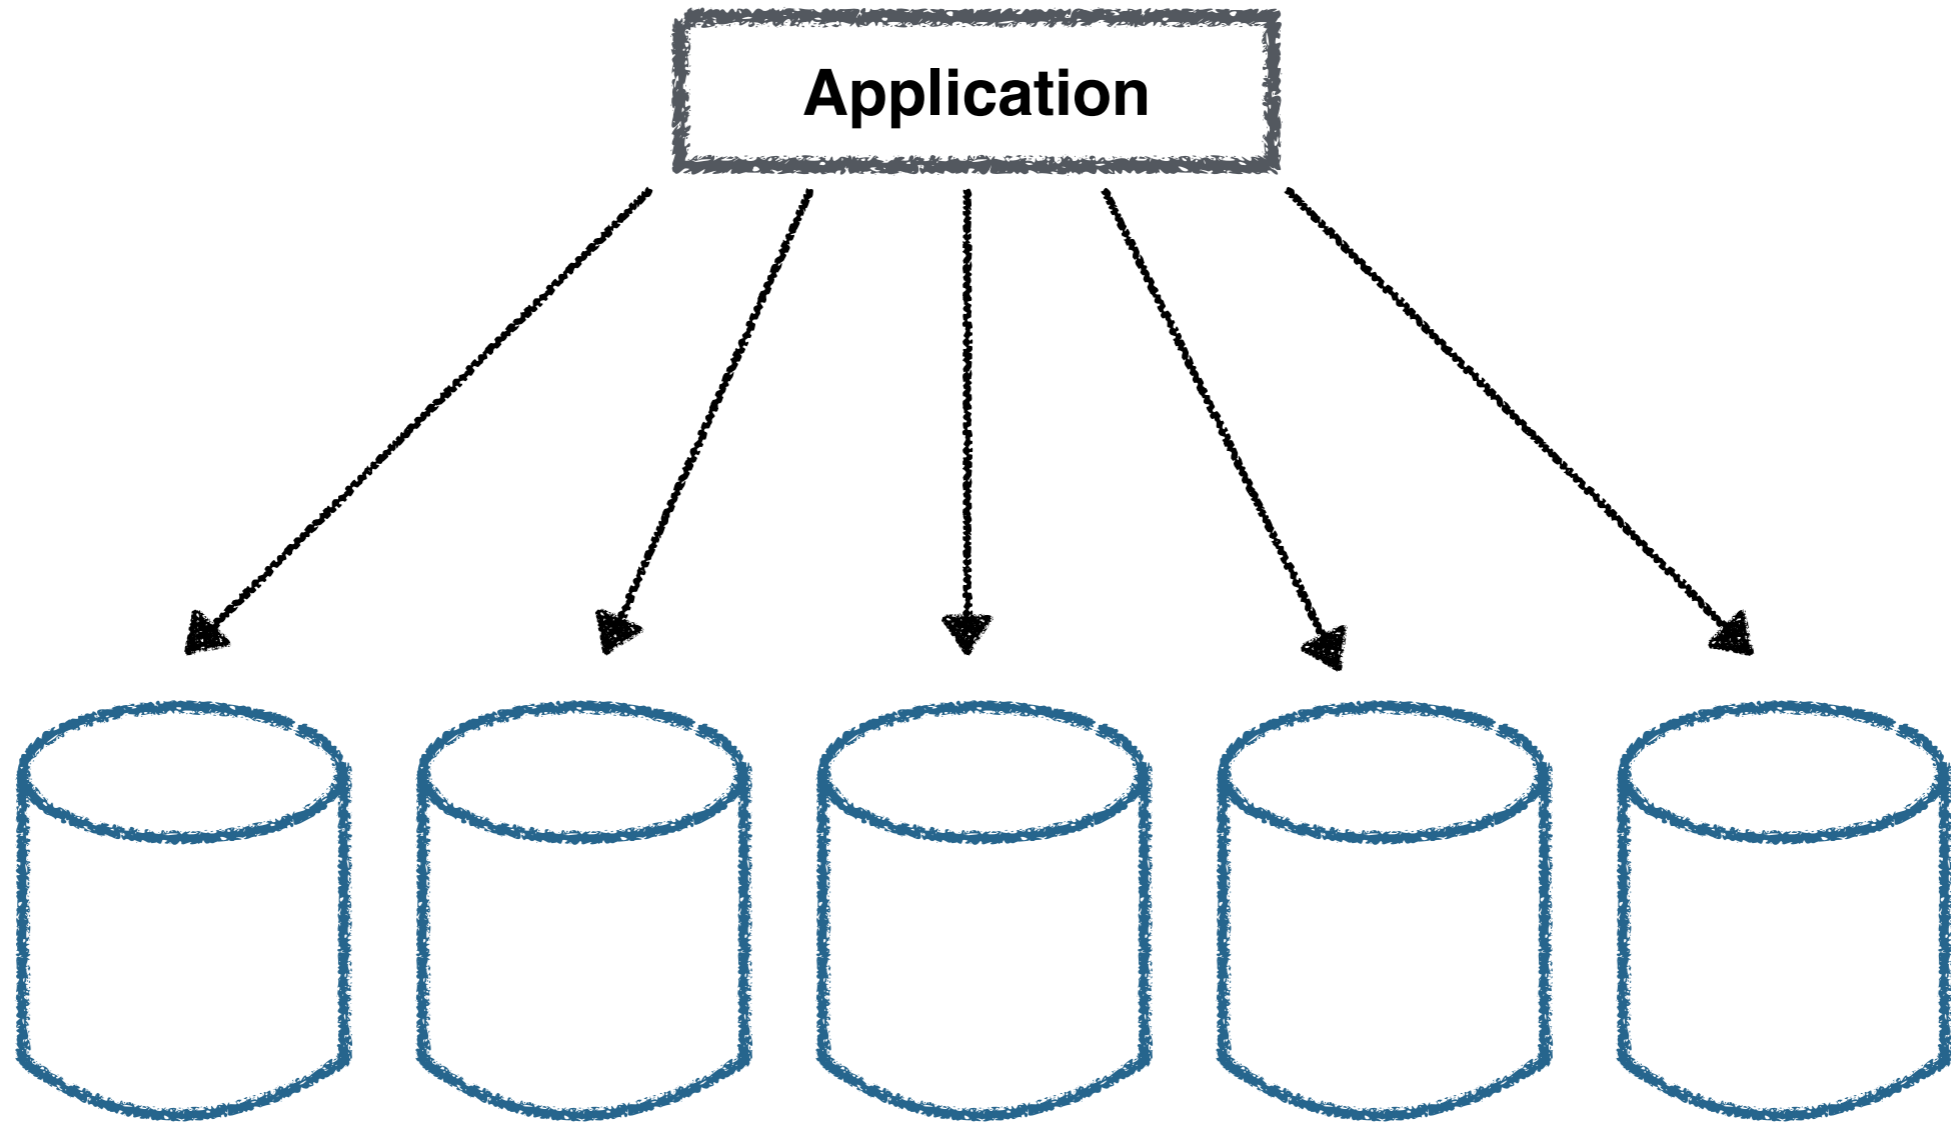
Transaction Manager / Planner



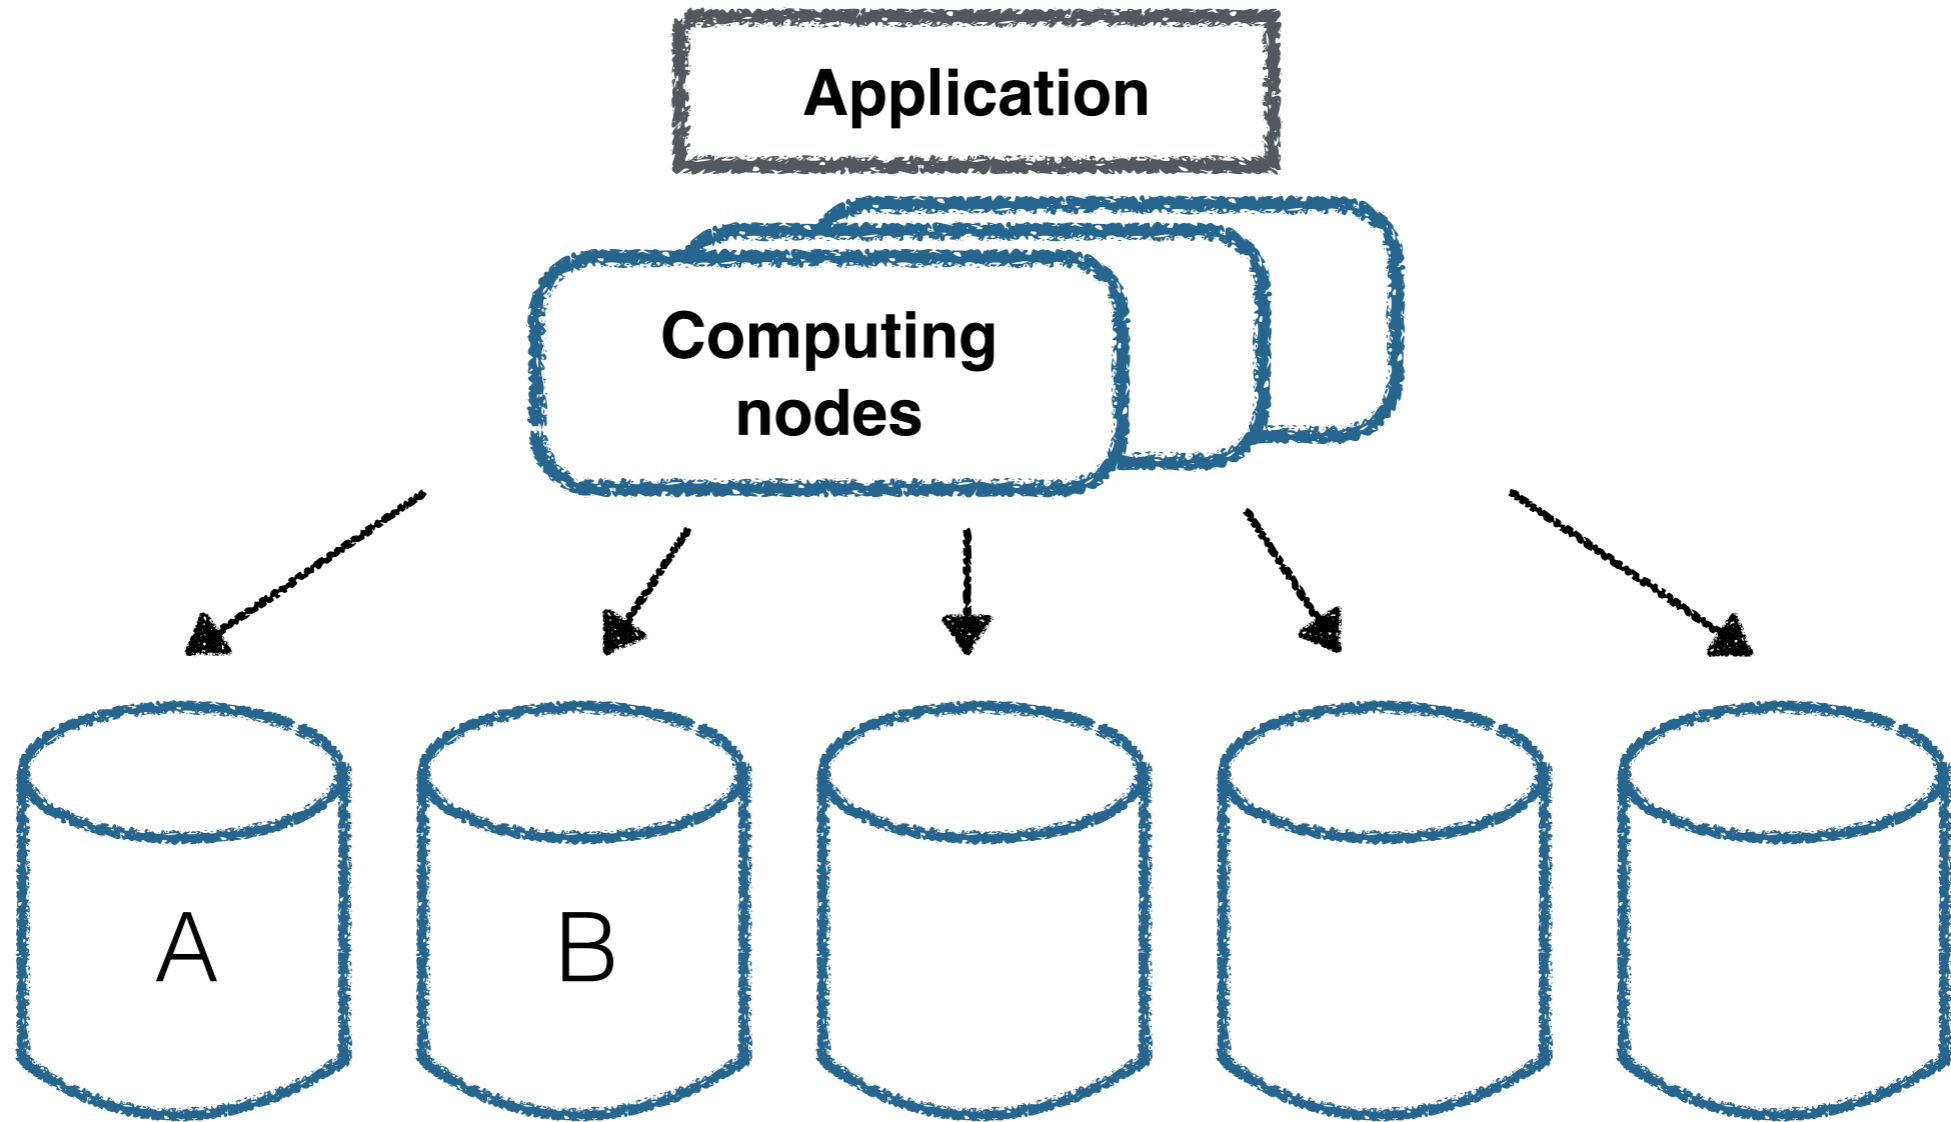
HOW TO WRITE CV



NOSQL Style



Map Reduce Style



Data Distribution


- Replicated
- Distributed by key
- Distributed evenly

Application

```
SELECT COUNT(*)  
FROM ...  
GROUP BY ...
```

Computing
nodes

```
SELECT SUM(n)  
FROM shards  
GROUP BY ...
```



```
SELECT COUNT(*) AS n  
FROM ...  
GROUP BY ...
```

Technology stack

- Pl/Proxy
- Pgouncer
- Skytools

Alternatives:

- PgPool
- Pg_shard
- Postgres XL
- Citus DB
- ...

Provisioning and monitoring

- Ansible
- Collectd
- Graphite
- Grafana
- Logstash
- ElasticSearch
- Redis
- Kibana

Alternatives:

- Puppet
- Chef
- Saltstack
- Zabbix
- Munin
- Nagios

Questions?

Have a nice day

sasha@shards.io