



How to test a storage system

watch your head not to explode!

Testing a storage system is....



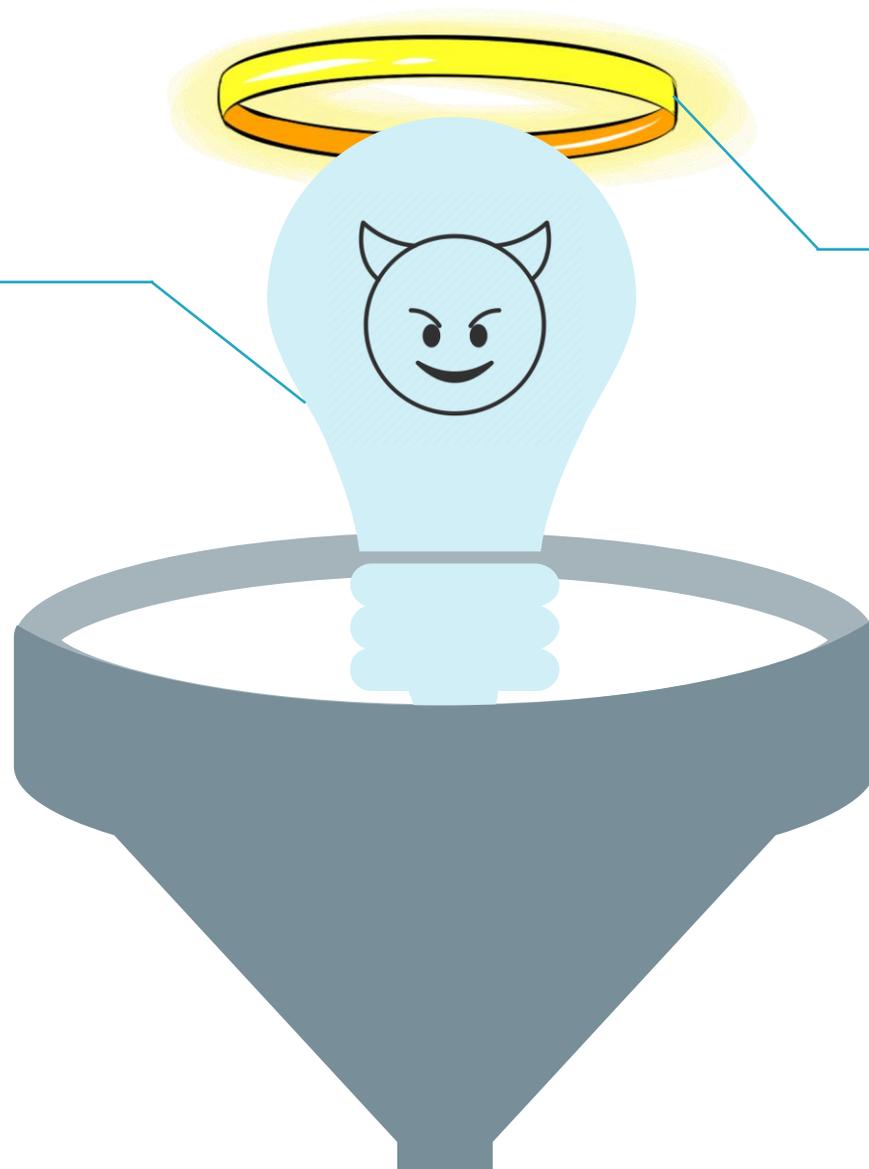
Testing a storage system is... **Versatile?!**

Versatile

Able to adapt or be adapted to many different functions or activities.

“Know-How”-driven

In fact, all Hi-end, mid-range and >50% cheaper storages include patented solutions

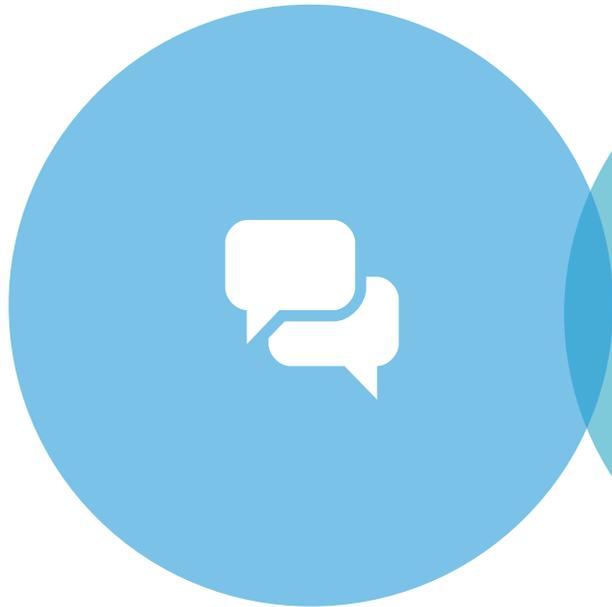


Storage system in a **nutshell**

**file storage
&
synchronization service**



I Storage Types



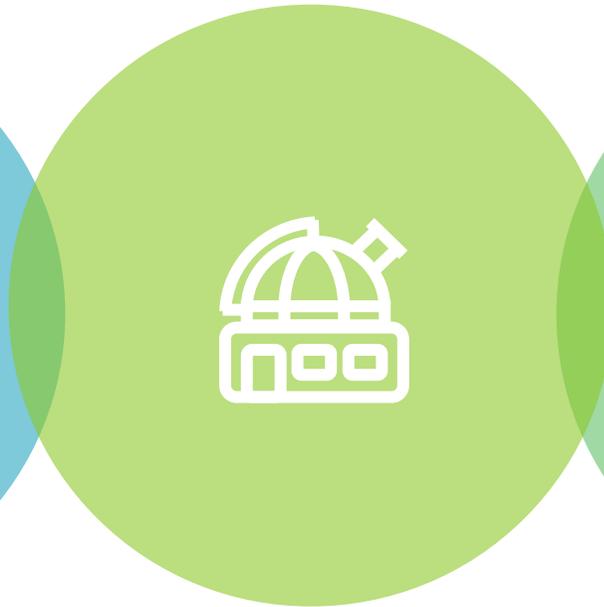
+ High Performant

- Social Networks
- Streaming Video



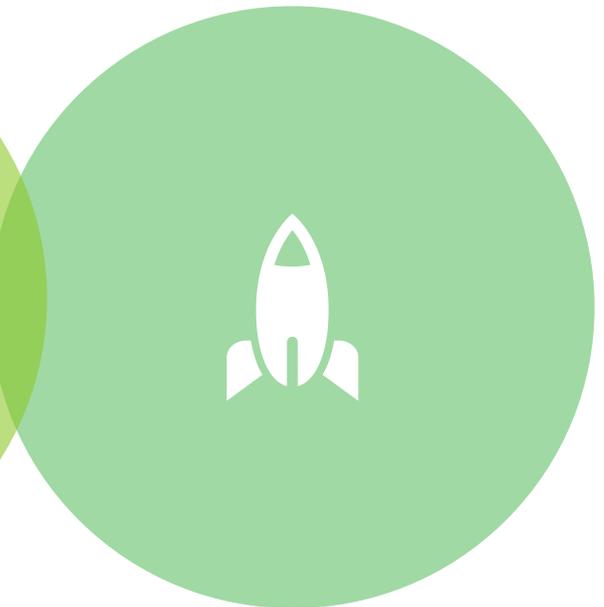
+ High reliability

- Banks



+ Fault tolerance

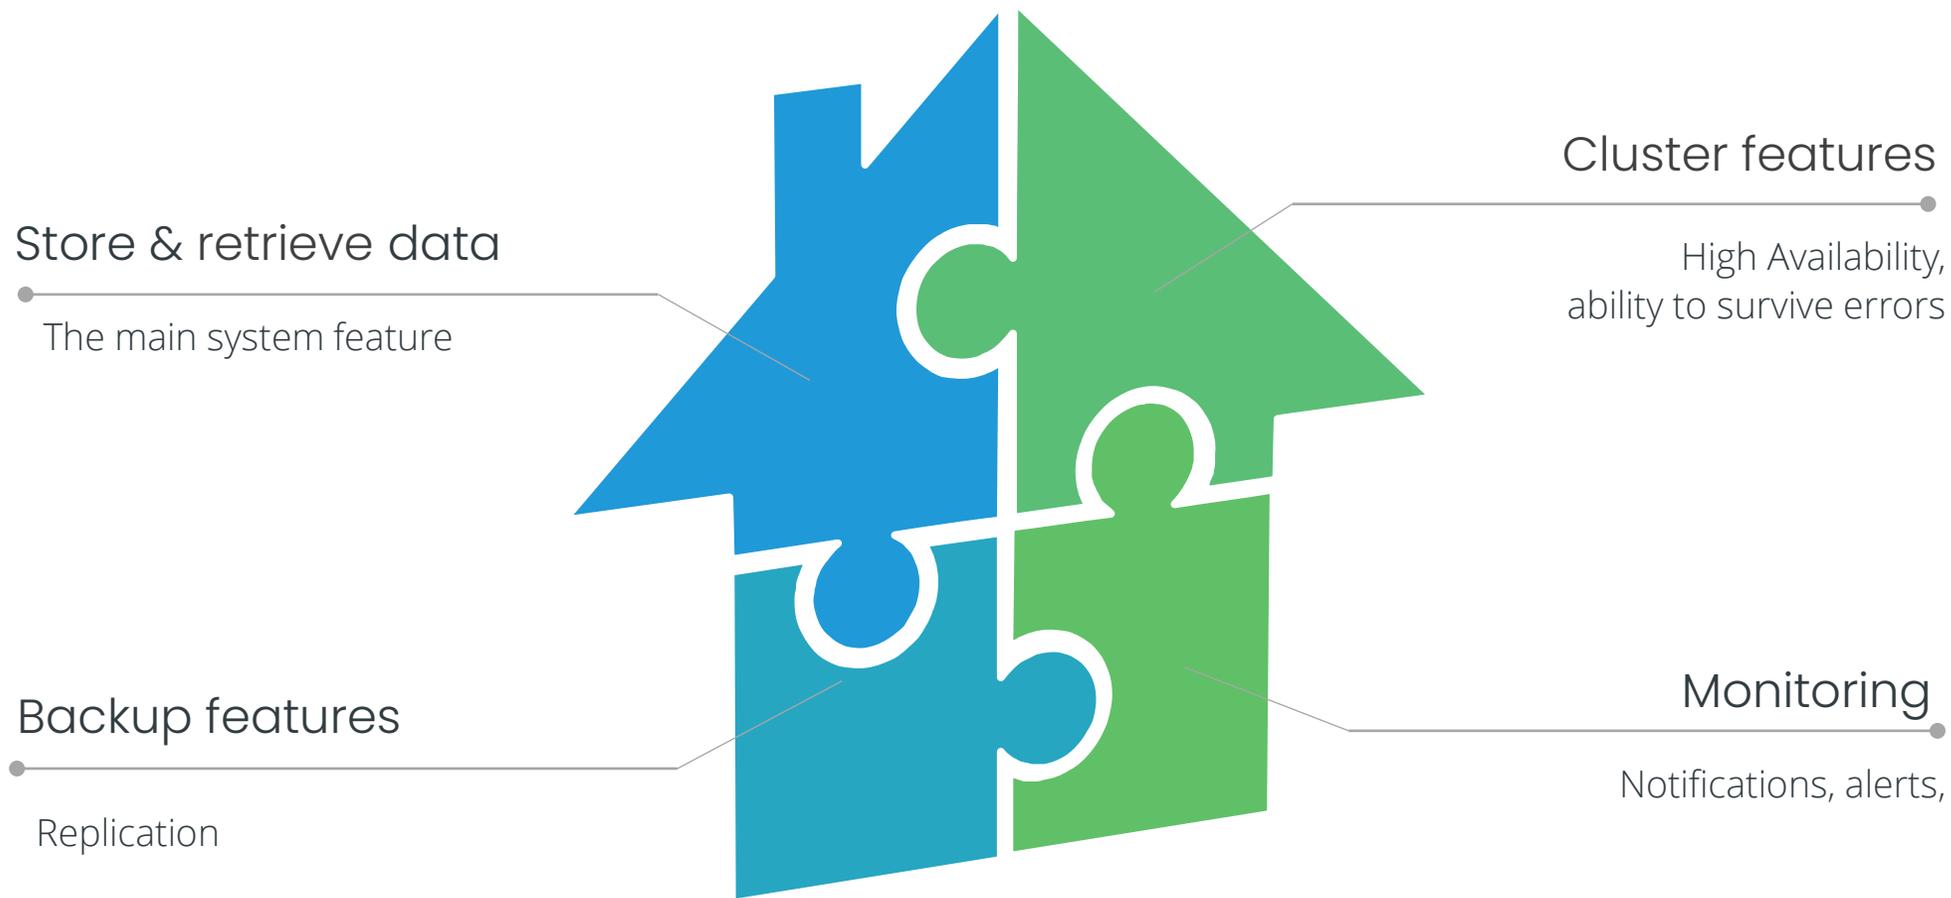
- Observatory



+ Hyper-converged

- Any business right now

I “Generic” storage system features

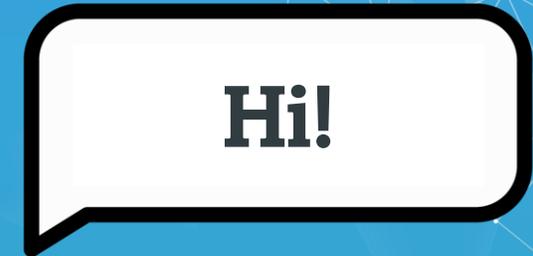
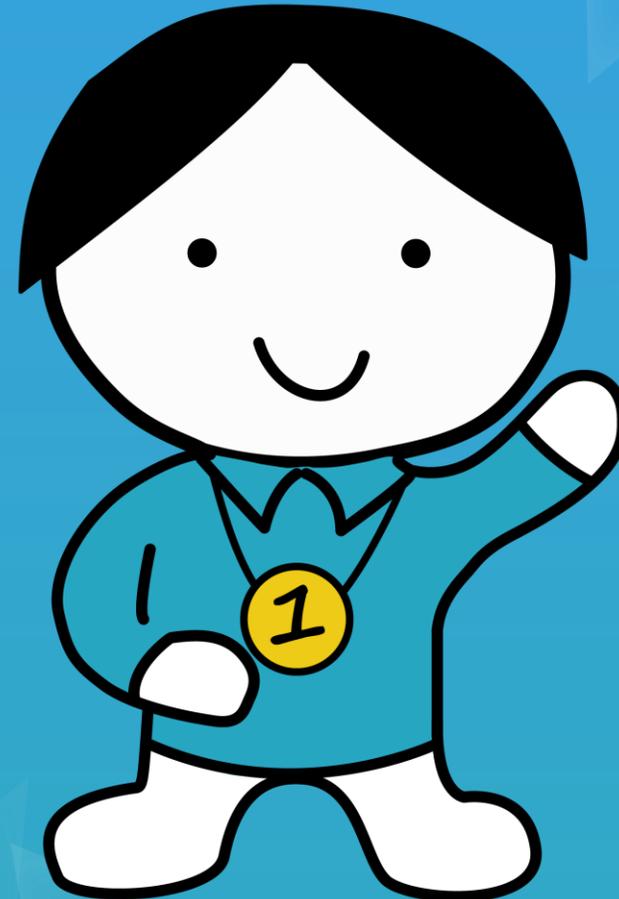




Storage QA engineer

Vasya Pupkin

Storage developing company



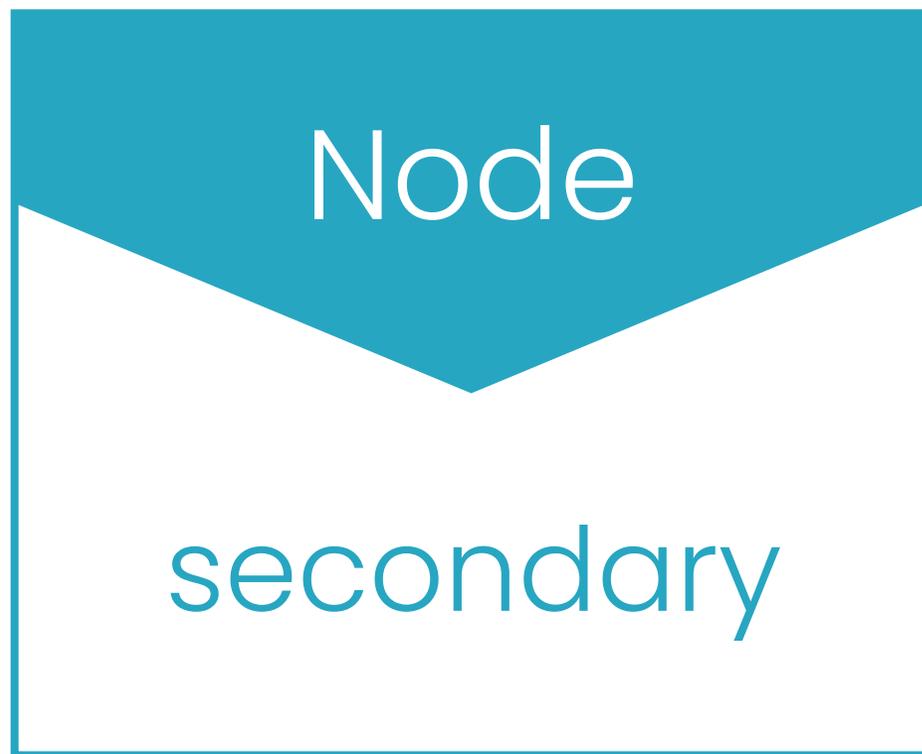
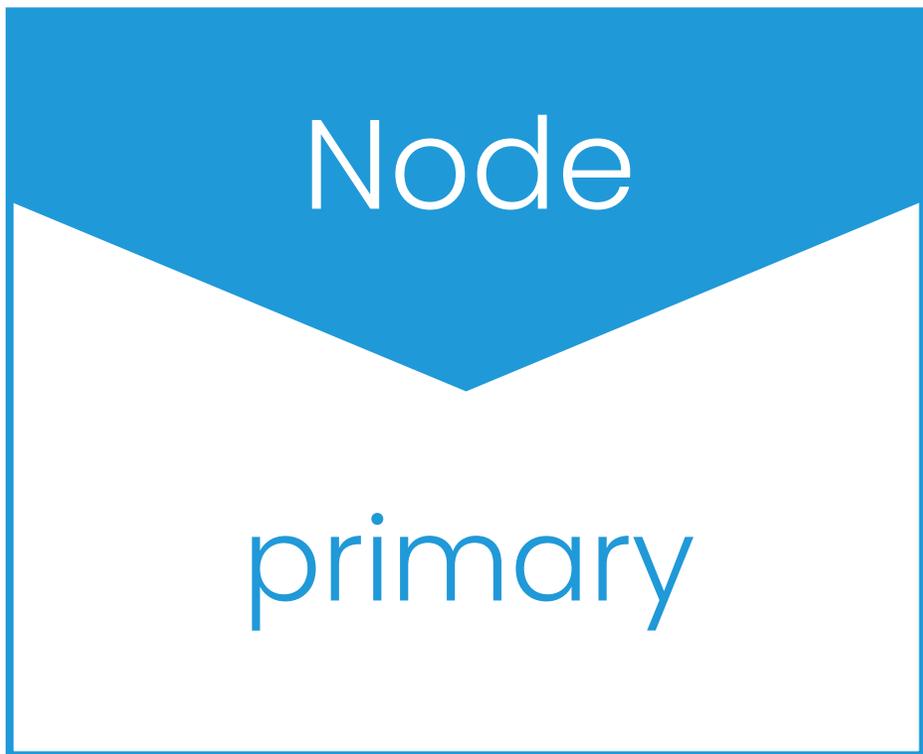
ALL CHARACTERS AND
EVENTS IN THIS SHOW--
EVEN THOSE BASED ON REAL
PEOPLE--ARE ENTIRELY FICTIONAL.
ALL CELEBRITY VOICES ARE
IMPERSONATED.....POORLY. THE
FOLLOWING PROGRAM CONTAINS
COARSE LANGUAGE AND DUE TO
ITS CONTENT IT SHOULD NOT BE
VIEWED BY ANYONE ■



Test #1:

"Upgrade is safe, right?"

Storage Cluster



Upgrade



Cold upgrade

Upgrade with stopping cluster and all production processes.



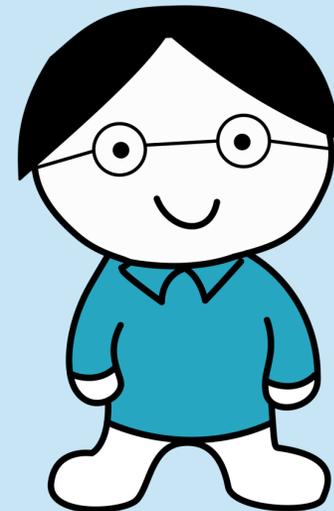
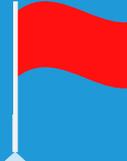
Non disruptive upgrade

Allows you to upgrade without production noticing it

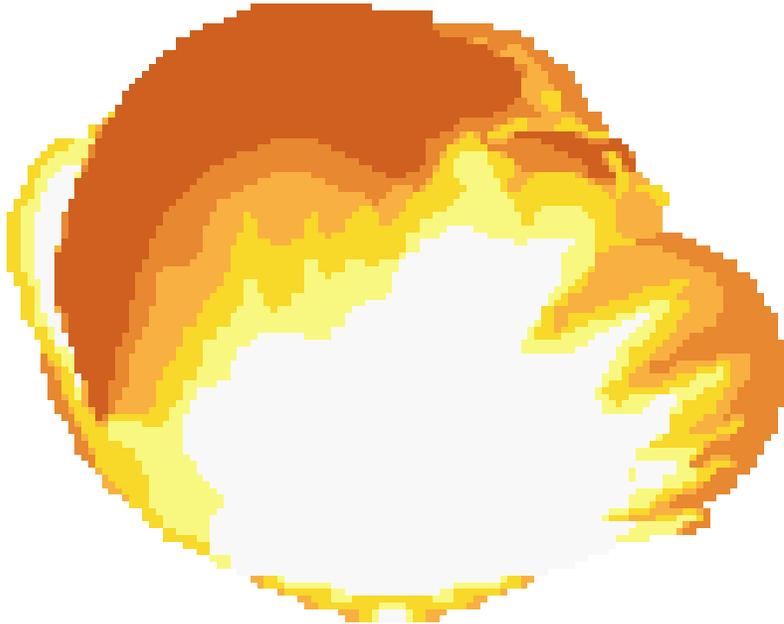
NDU to new version

Expected result: 2 nodes

Cluster version is N+1



Actual result:



Clues:



Data is gone = DU

Sad_cat.jpeg



Web UI is dead

“What means cluster is not operational!?”



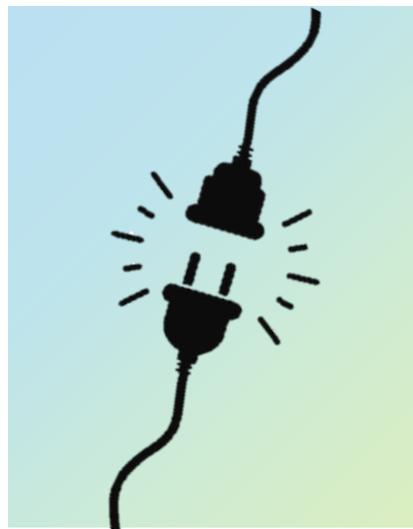
Cannot login to Nodes

No more SSH for you

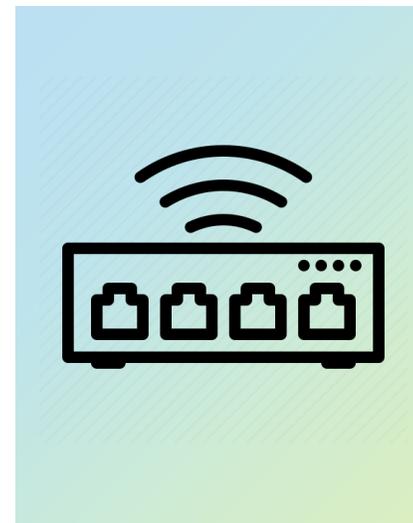
Hypotheses



Stolen by storage gnomes



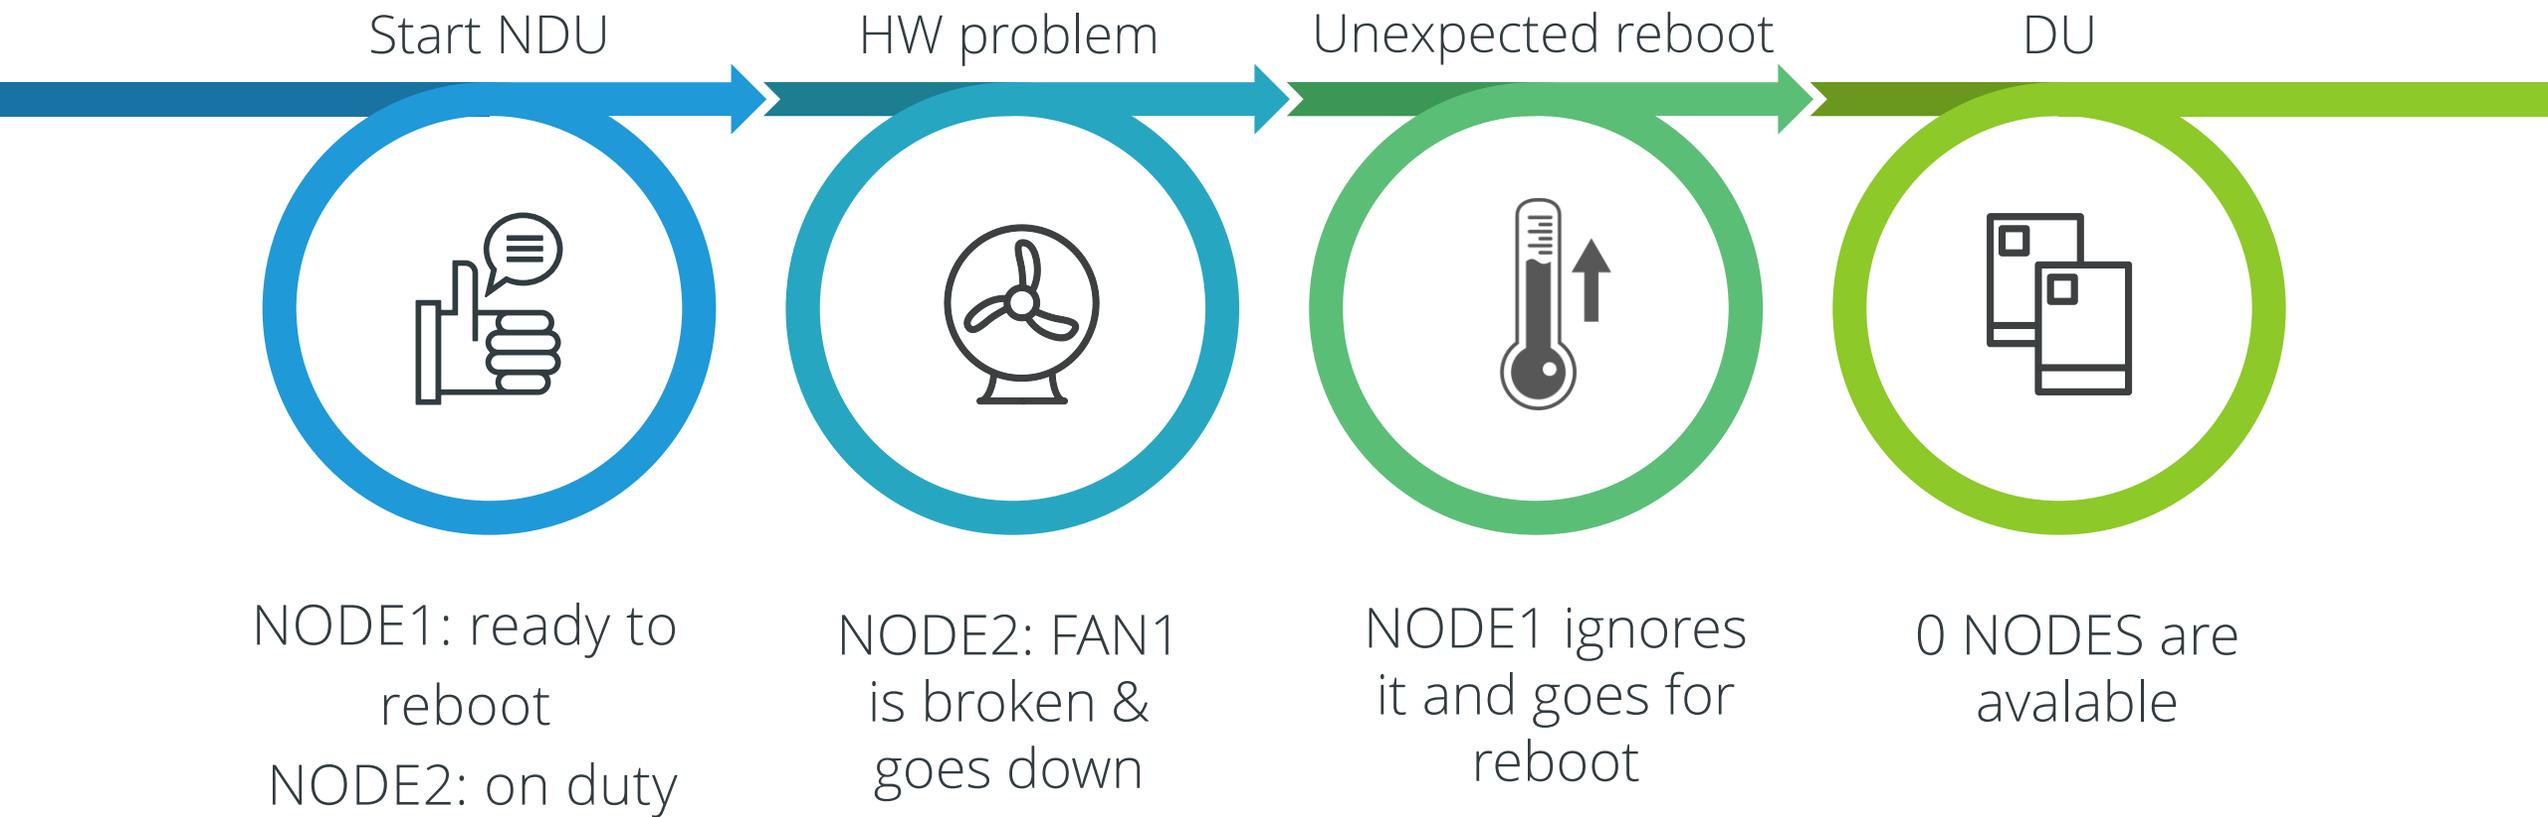
Physically disconnected



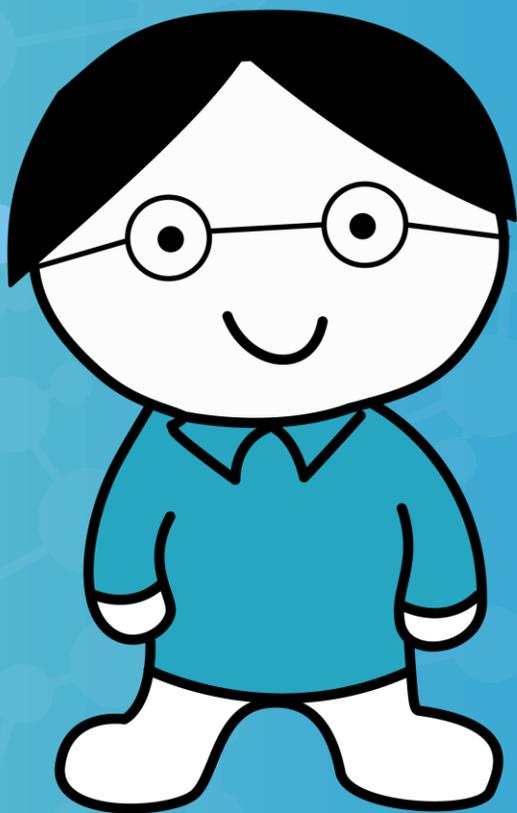
Hardware problem Just in a right time



BUG scenario:



The BUG



NODE1 decided to reboot despite it was
the last available cluster node



Test #2:

"Statistics? There is never enough!"

What are counters and events



Counters

increment if something happened
(counter += 1)

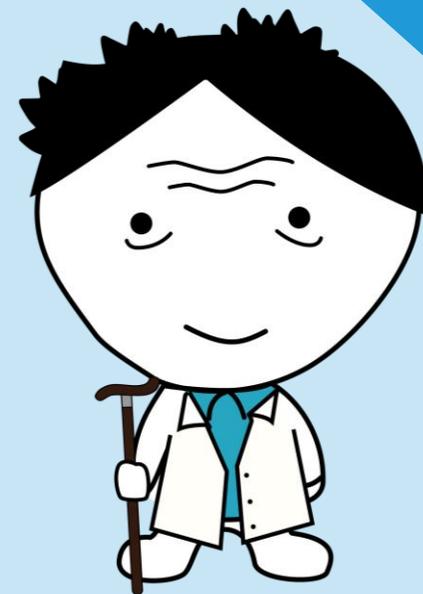


Events

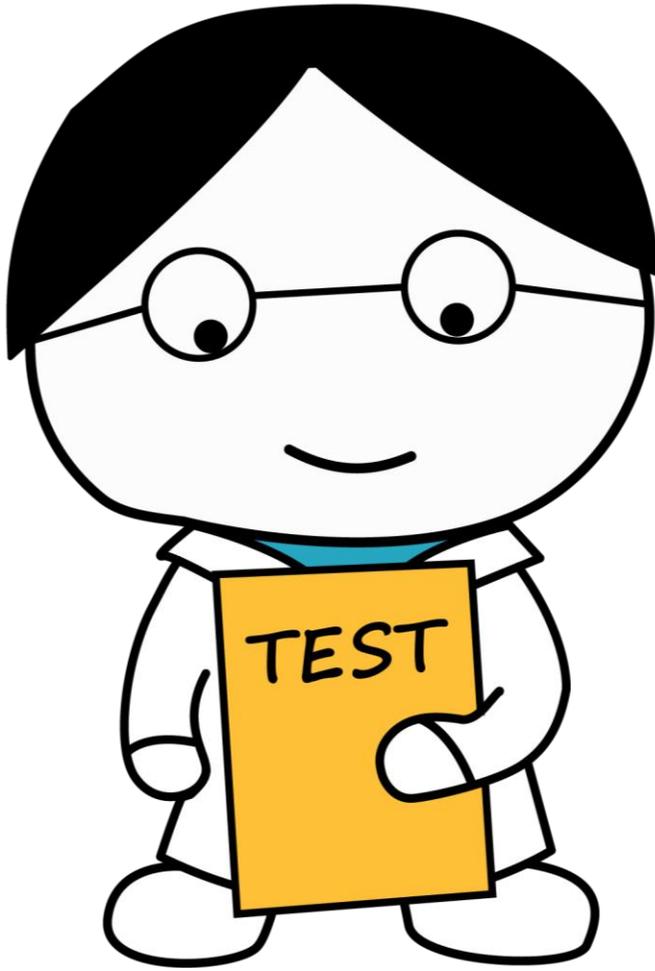
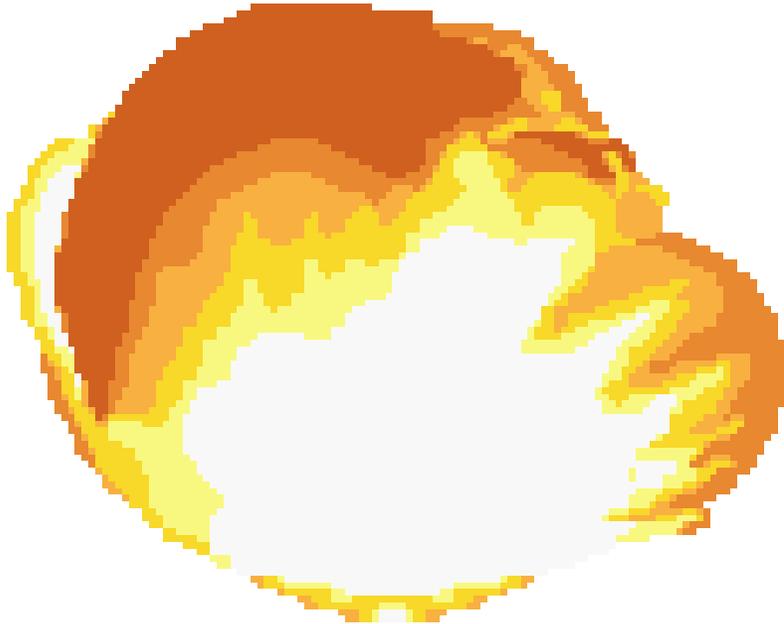
rise if something happened
(if counter == X)

Dump ALL stats every 1 sec by REST

Expect result: tasty stats are
collected, and all is fine



Actual result:



Clues:



Reads OK, Writes NO

Not so bad, this is not a DU, but 1\2 of main funcs are not working



MailBox is burning

100500 email notifications say "no more storage left in the system"



No more storage

3GB of 1000TB are used by UD
Where are 997TB?

Hypotheses



Stollen by storage gnomes



All Drives are empty



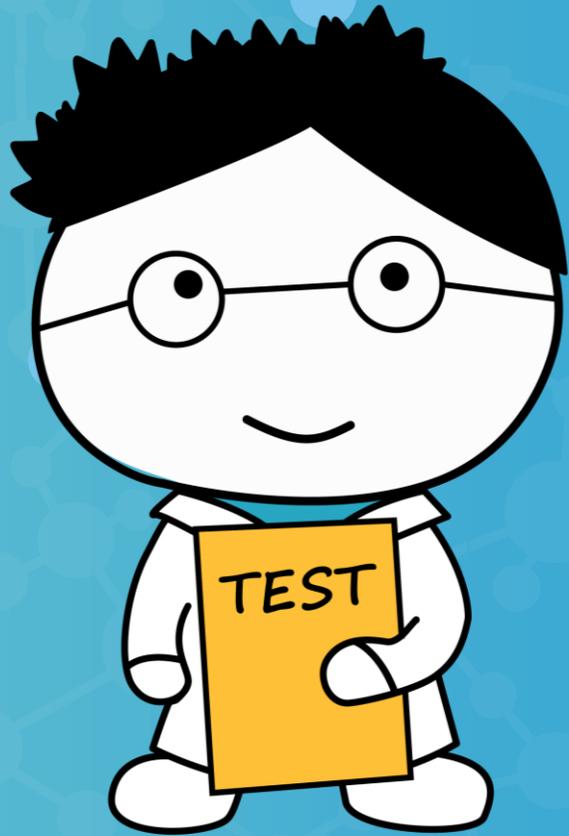
Disks are used by the system itself



The Cascade



The BUG



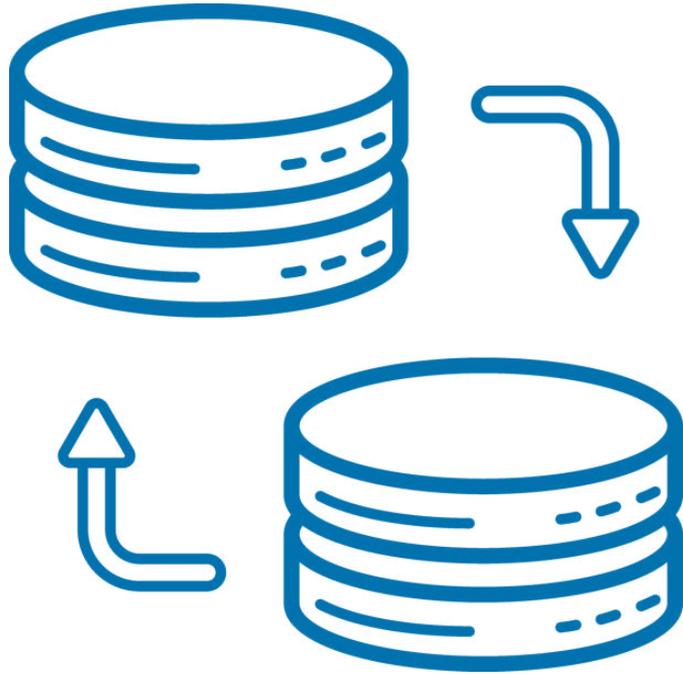
The MAX quantity of simultaneously kept REST-requested stats dumps is not limited and may lead to mgmnt DB over growing



Test #3:

"Wanna be safe - do backups!"

Replication and RPO (in short)



Replication

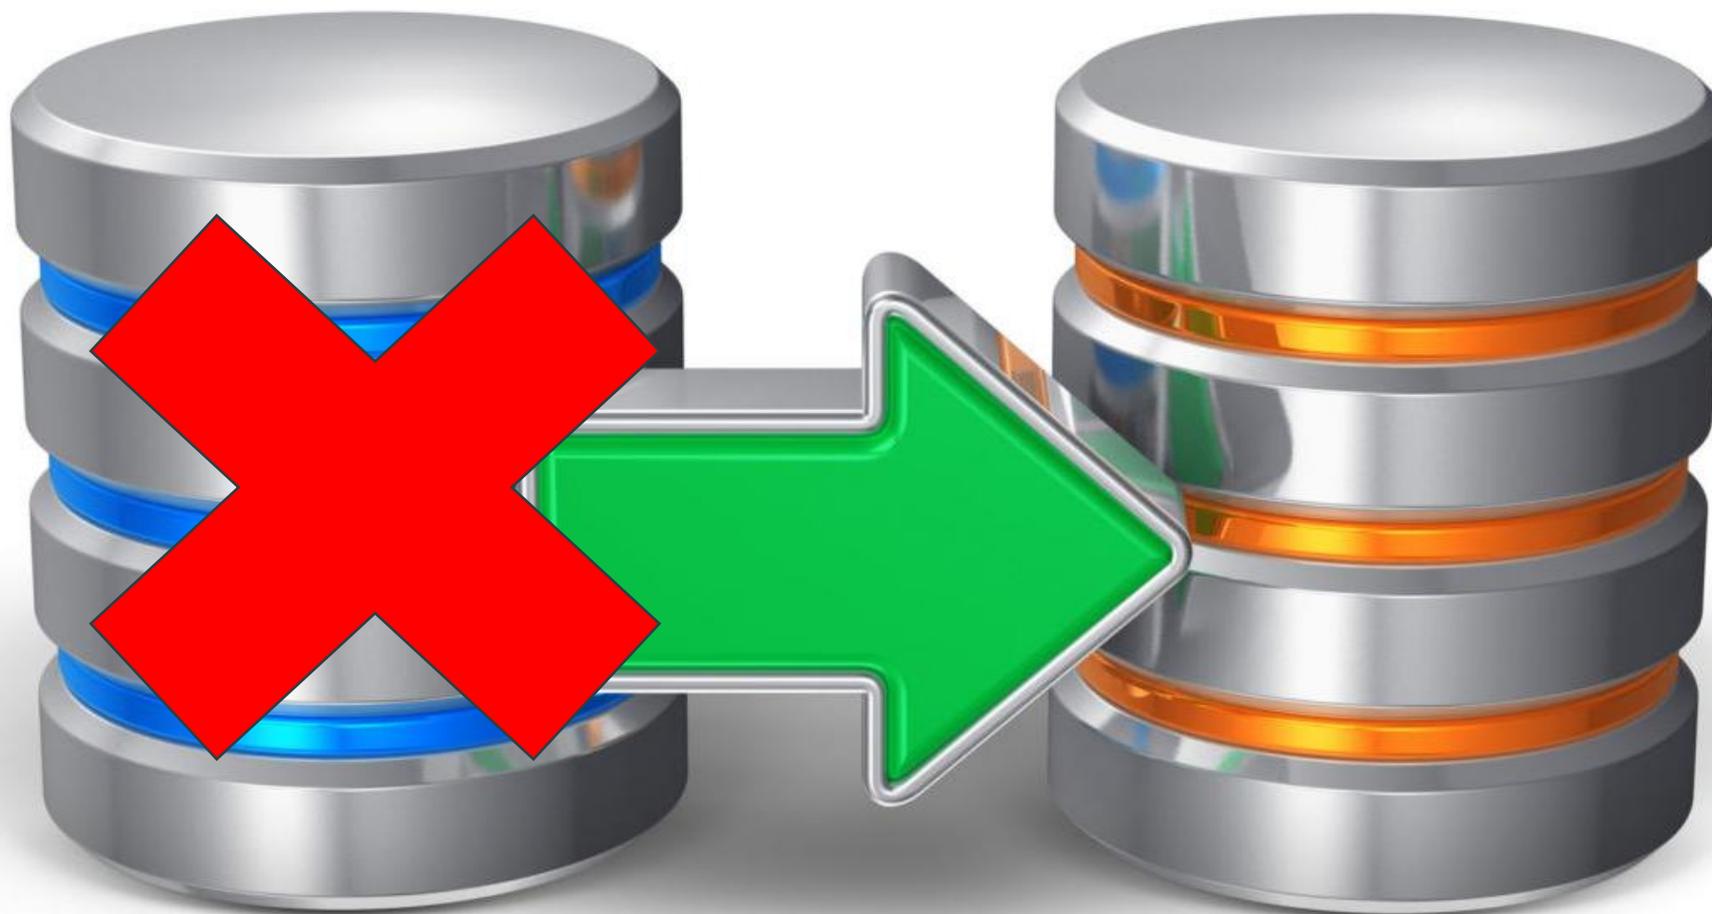
sync-ed copies of data are distributed between sites



RPO

Recovery Point Objective
(how many info we may lose,
measured in min)

Typically made between 2 geographical sites



Tokyo --> NY

Replication

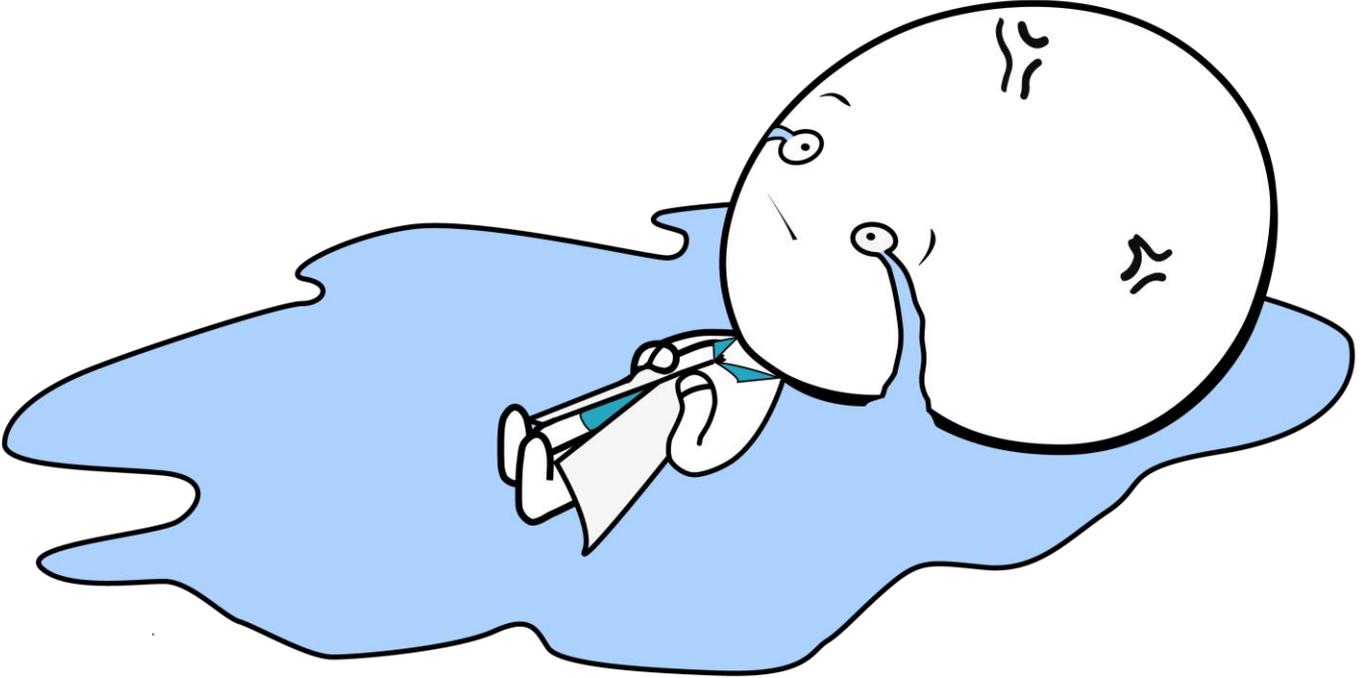
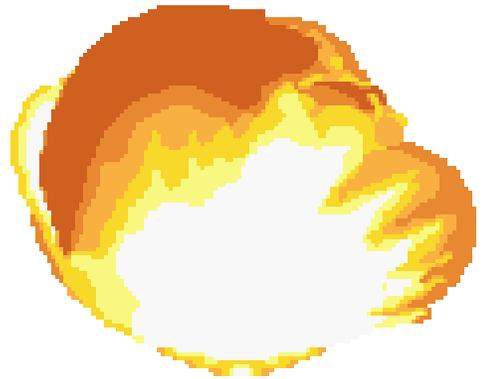
RPO 5 min

Expect result: Storage is
secured, RPO completed



PROF. VASYA X

Actual result:



Clues:



Web UI is dead

Again



Write\Read is ok

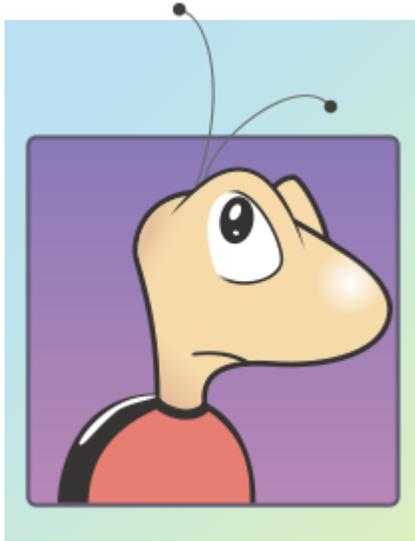
Clients are in Lab A, access the storage by wires



Labs are gone

Tokio and NY Labs don't respond to ping...

Hypotheses



The Bugzilla



The storage gnomes

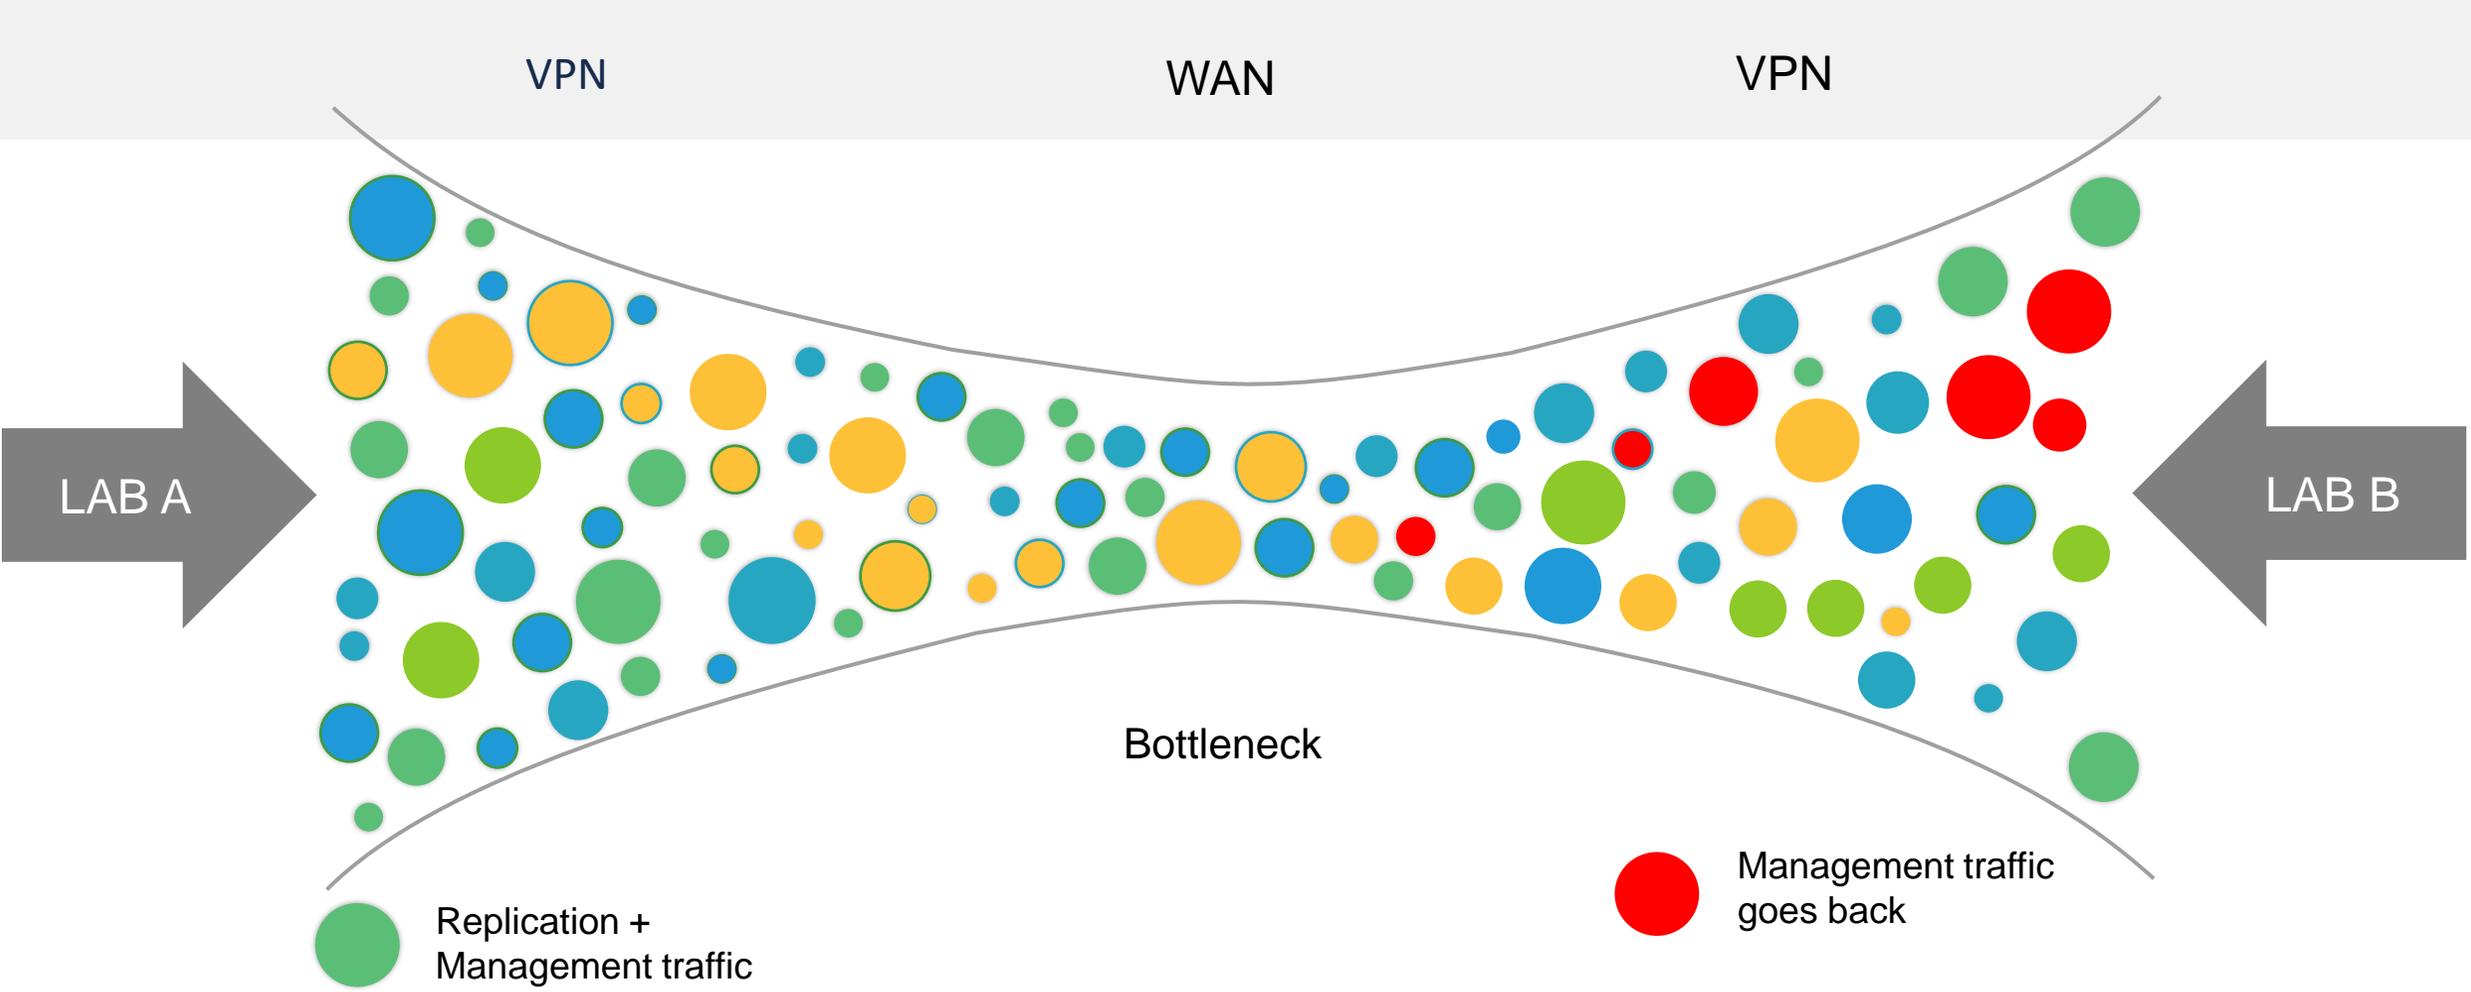


Connection



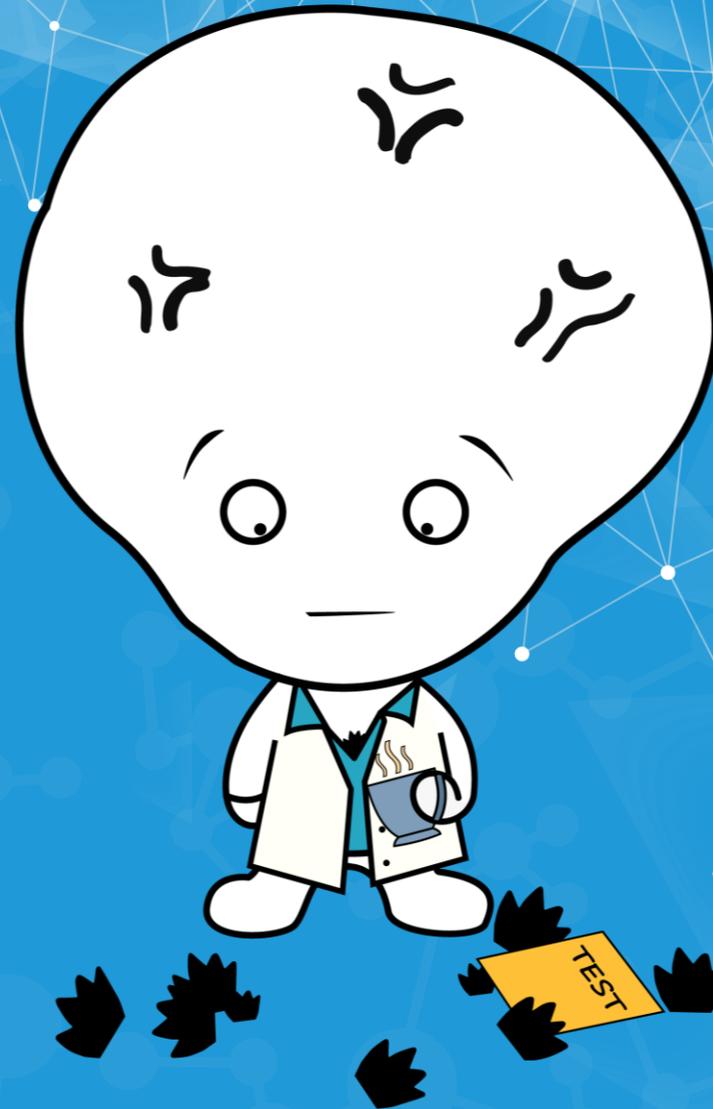


The overflow



The BUG

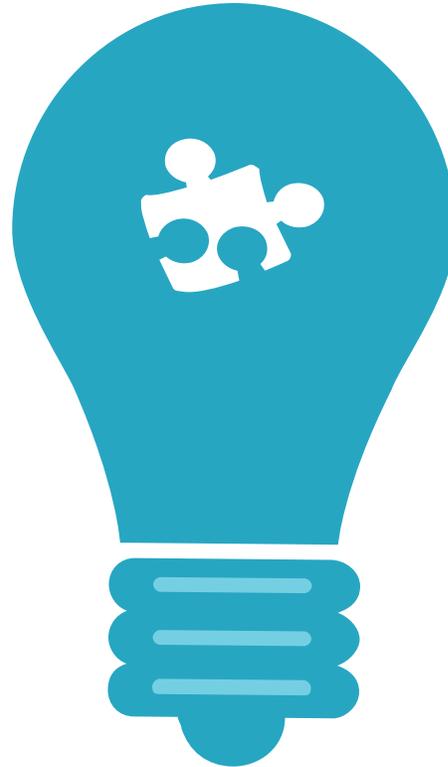
Replication has no QoS policies for management & data traffic balancing



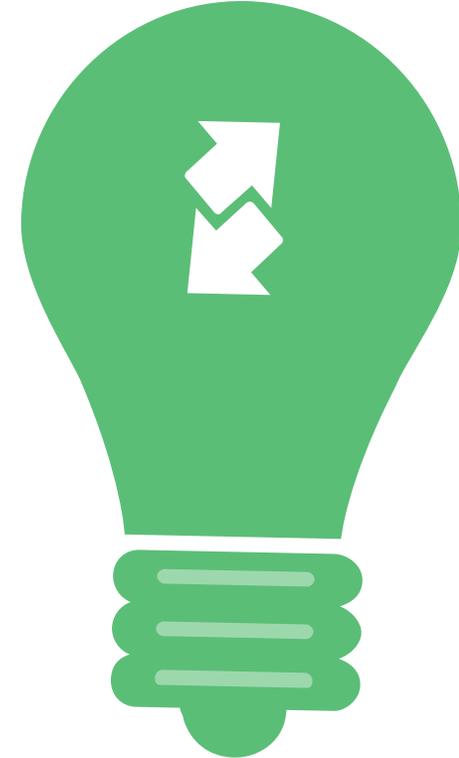
| How to test a storage system?



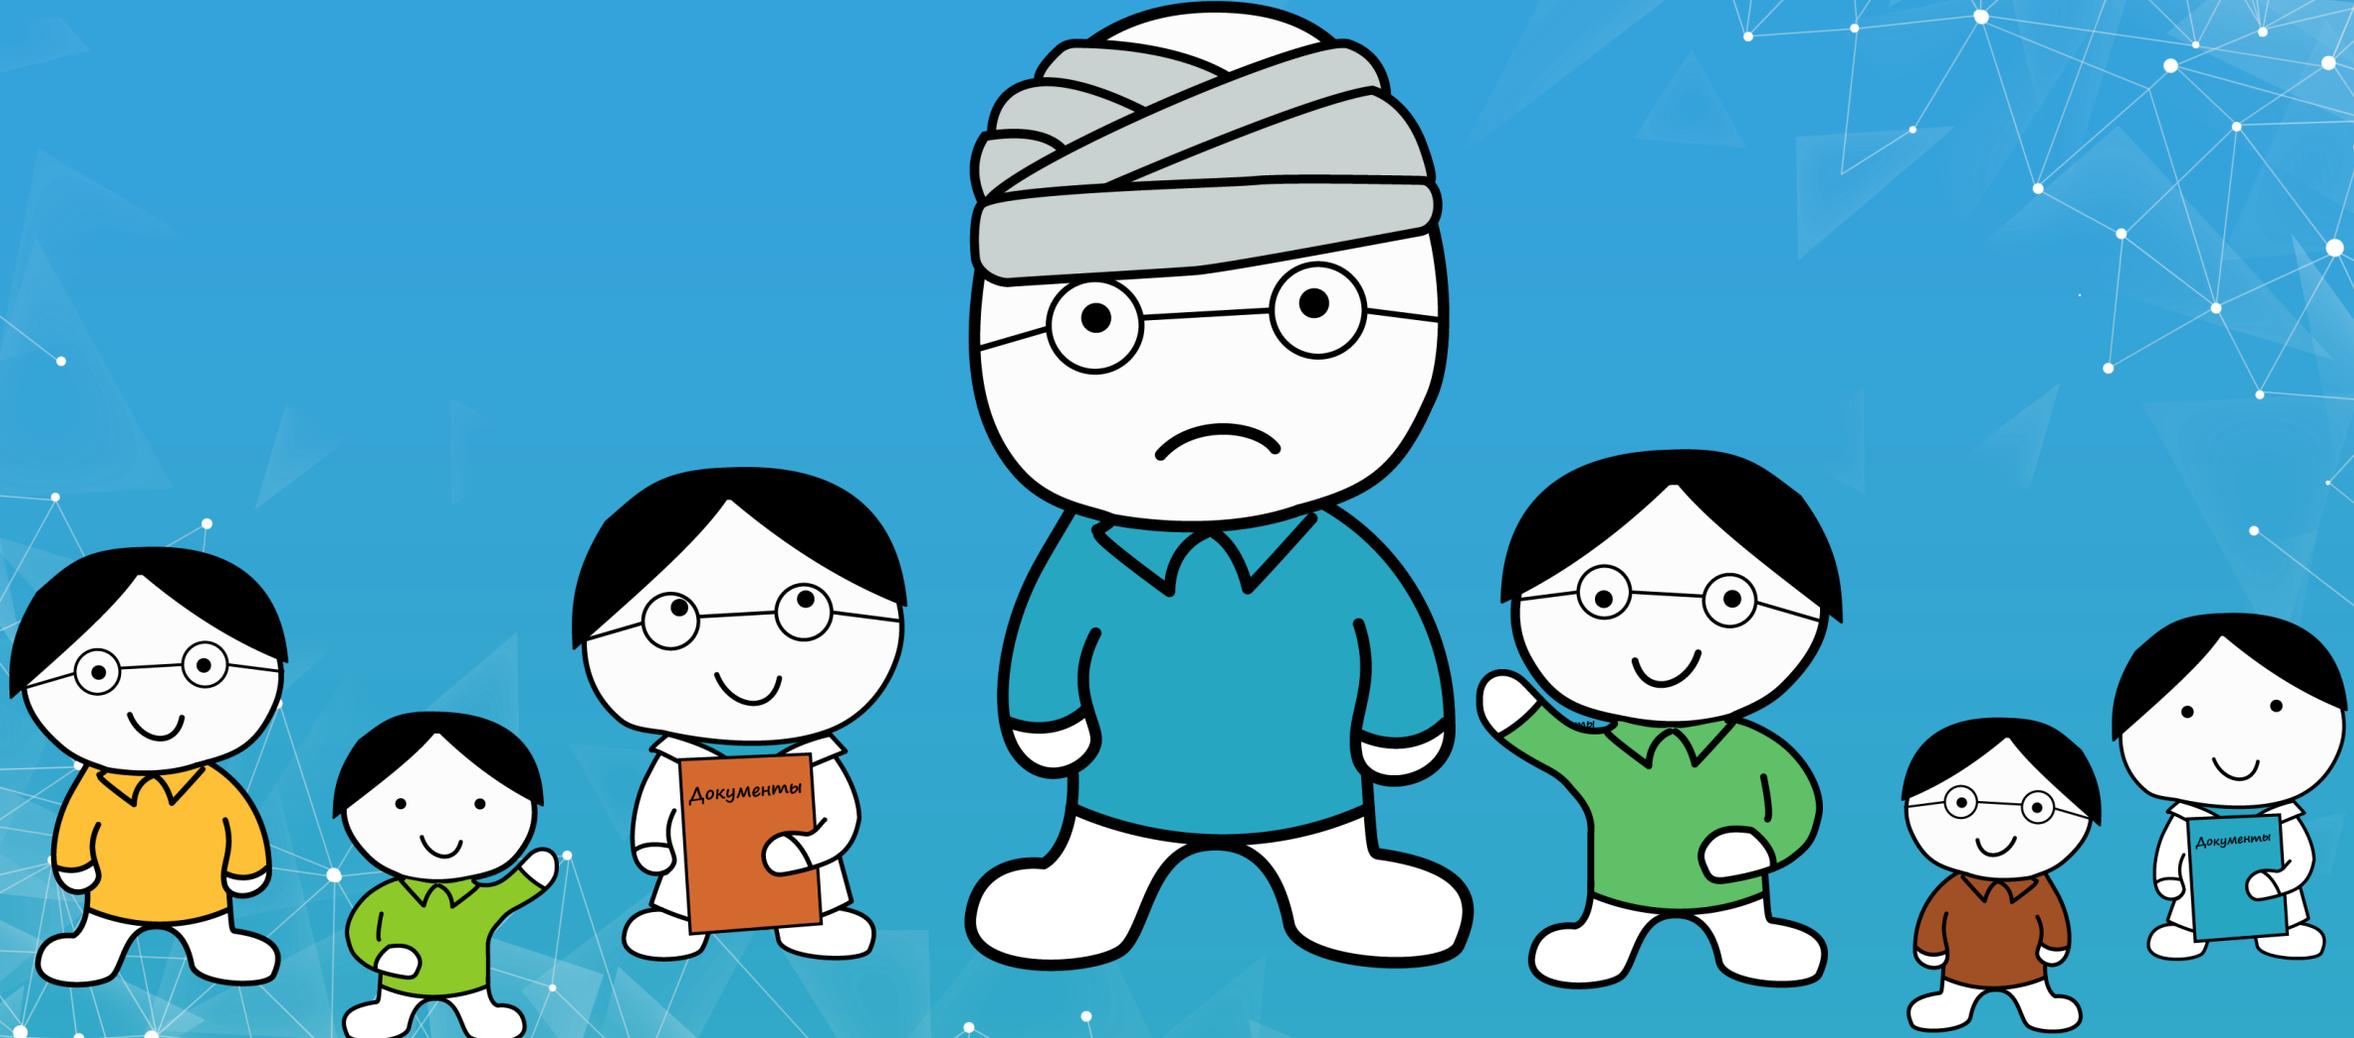
**Focus on interesting
Knowledge Domains**



**Be a team player,
create own network**



**Beware of
STORAGE GNOMES**





Q&A session

Thanks for watching