Migration Math

Thorsten Lange
Storage Architect
Sun Microsystems

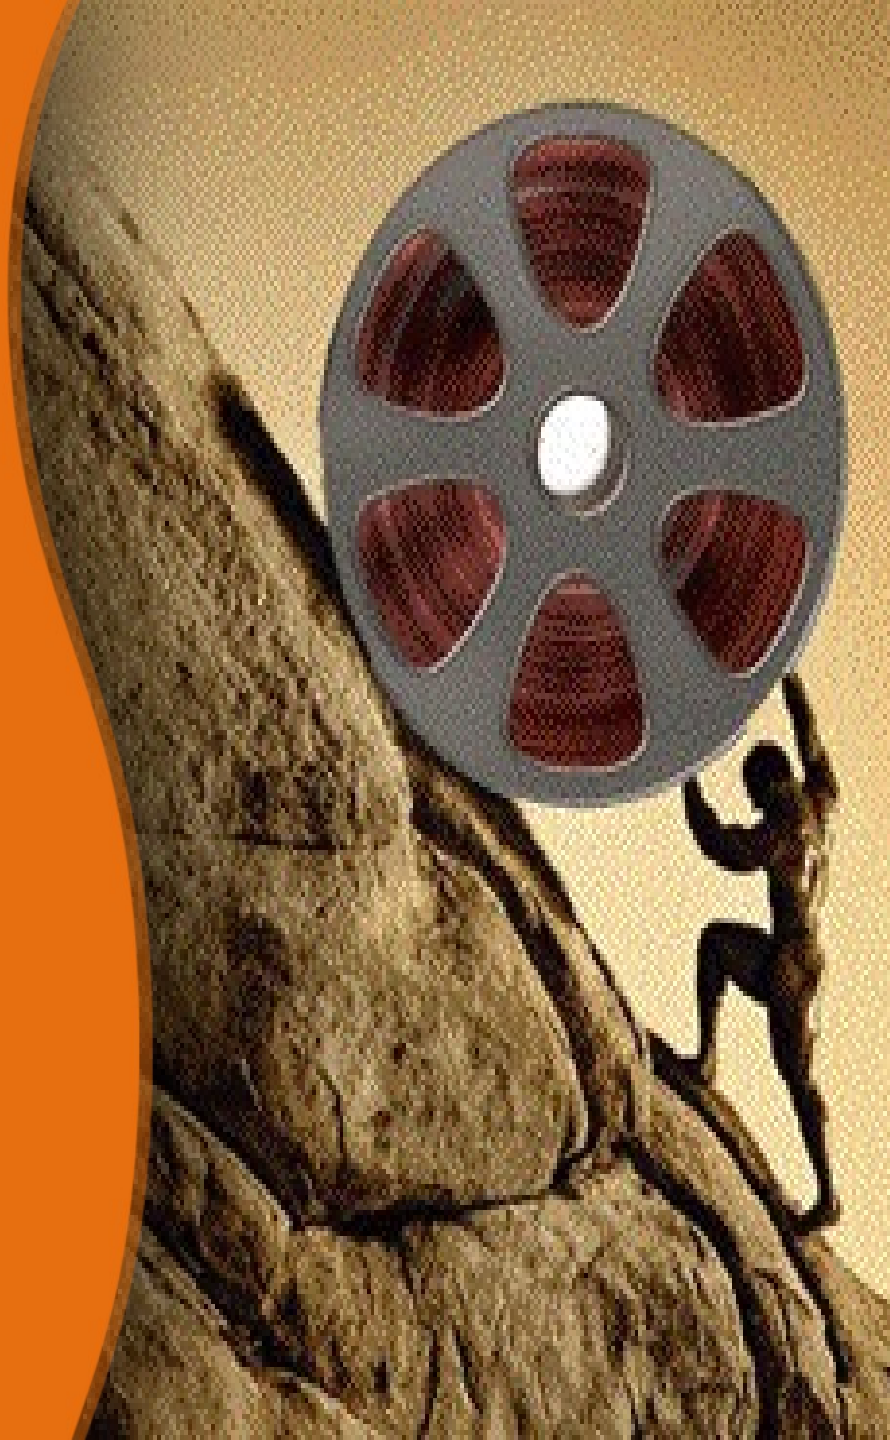

Q: How do you preserve boiling water?

Some suggestions

- Changing the water is not an option
- Leave the flame on for the next 100 years
- Put it in the fridge
  - Sue the manufacturer because the water is no longer boiling when you take it out
- By the way – what exactly is the re-use case?

Some suggestions

An appliance is the key !!!

Separate content from condition

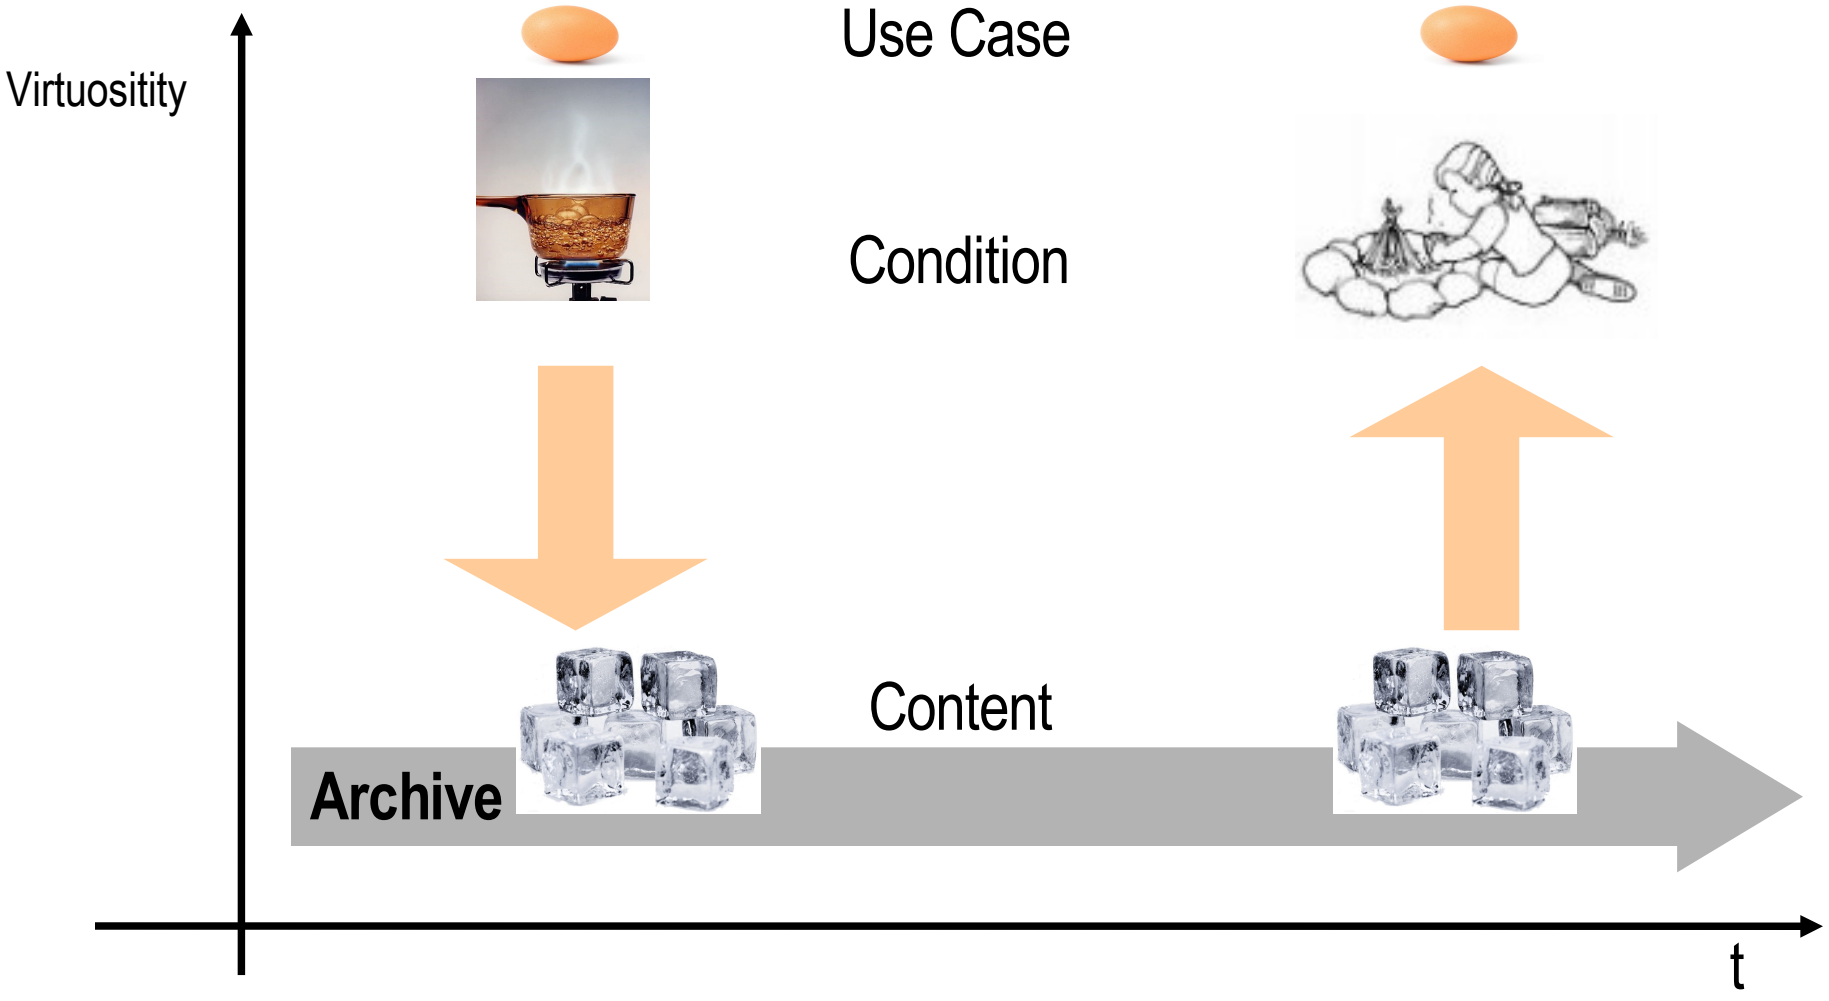

Re-Use in the Far Future?

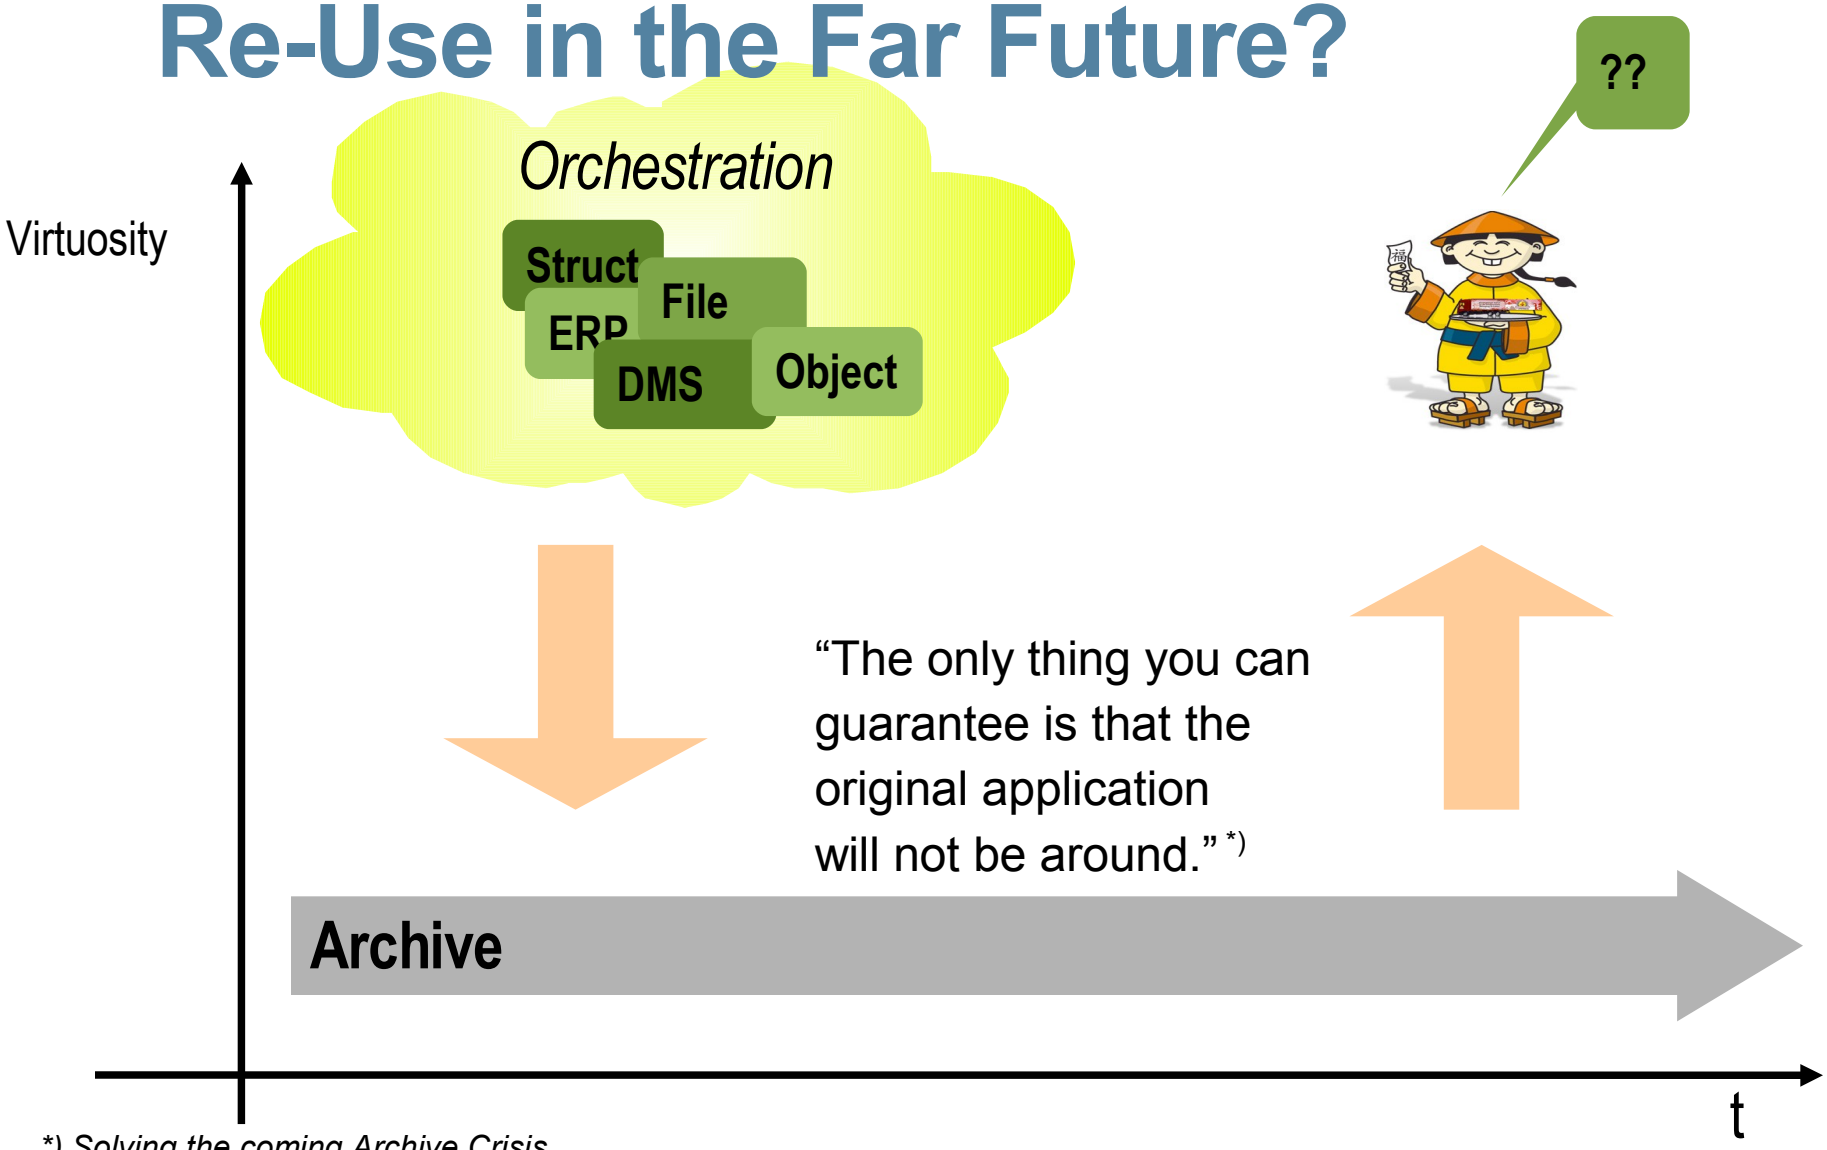

*) Solving the coming Archive Crisis
http://www.snia.org/forums/dmf/programs/ltacsi/100_year/

The order of an algorithm

- Remember when you learned to code an algorithm?
- Let's take a simple example:
  - I give you a list of words, and your task is to identify which of these are anagrams.
  - Anagrams are words formed by the same letters, just resorted. Like:
    - MARY – ARMY
    - GEORGE BUSH – HE BUGS GORE

The order of an algorithm

- 1st idea:
  - Take the first word an determine all its anagrams (function *permute*). Compare all these with the other words and print it in case of a match. Then continue with the 2nd word etc.
  - Feeding this program with 10 words of 4 letters, it's back in 1s (I'll assume a computer speed of 100 actions per second)
  - Feeding WHATAFUNNYPROBLEM into *permute* takes a while ...
    - $3.5 * 10^{14}$ permutations
    - 111,243 years ...
  - The order of my algorithm is
    - O(n!)

The order of an algorithm

- 2nd idea:
  - First sort the letters of each word in its alphabetic order. Then compare.
    - O(n*log(n))
  - If a 10 letter word takes 1s, I expect a 1,000 letter word be around 300s, and so on.

Some assumptions

- What I experience today is about the amount of data that I'll need to handle for all times
- The overall amount of archived data will shrink
- Future research will result in fewer data that requires archiving
- Technological development will continuously reduce my effort to migrate my stuff into the future
- People will find out that a lot of that archived material isn't worth to be kept any longer, so I can shredder it in time

Are you with me ... ?

Data Growth

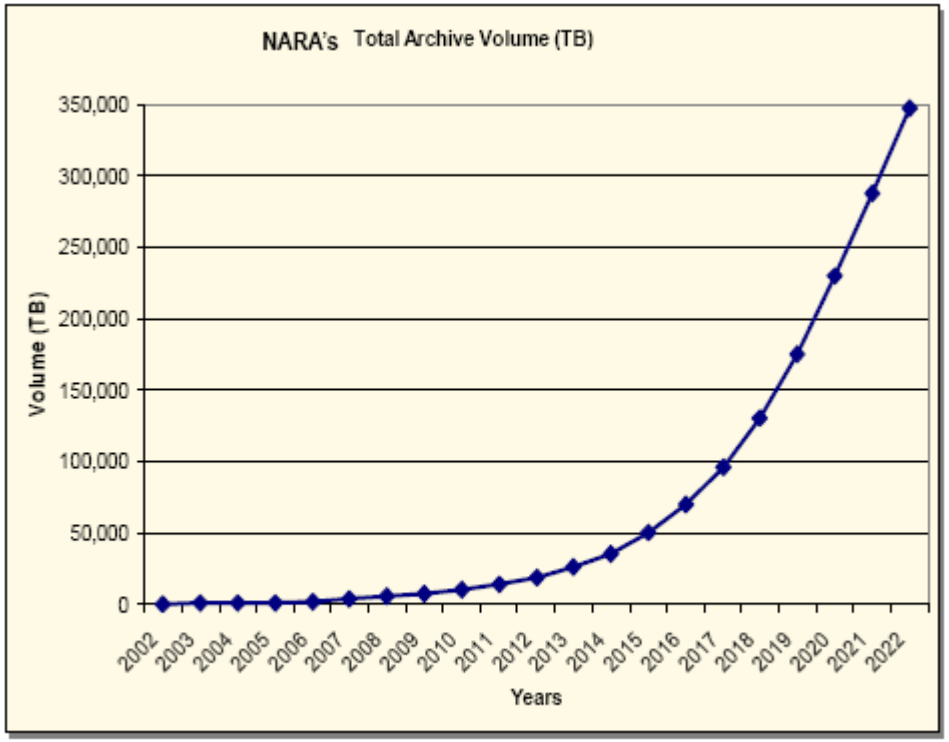

Source: Ken Thibodeau, Director ERA

The Migration Challenge

An organization that has 1 petabyte, (PB) in its digital archive repository will have 50% more next year. In three years, they will need to migrate that first petabyte. In five years they will need to migrate 2.25 PB.

100 Year Archive Requirements Survey

Solving the coming Archive Crisis

http://www.snia.org/forums/dmf/programs/ltacsi/100_year/

Principle of Migration

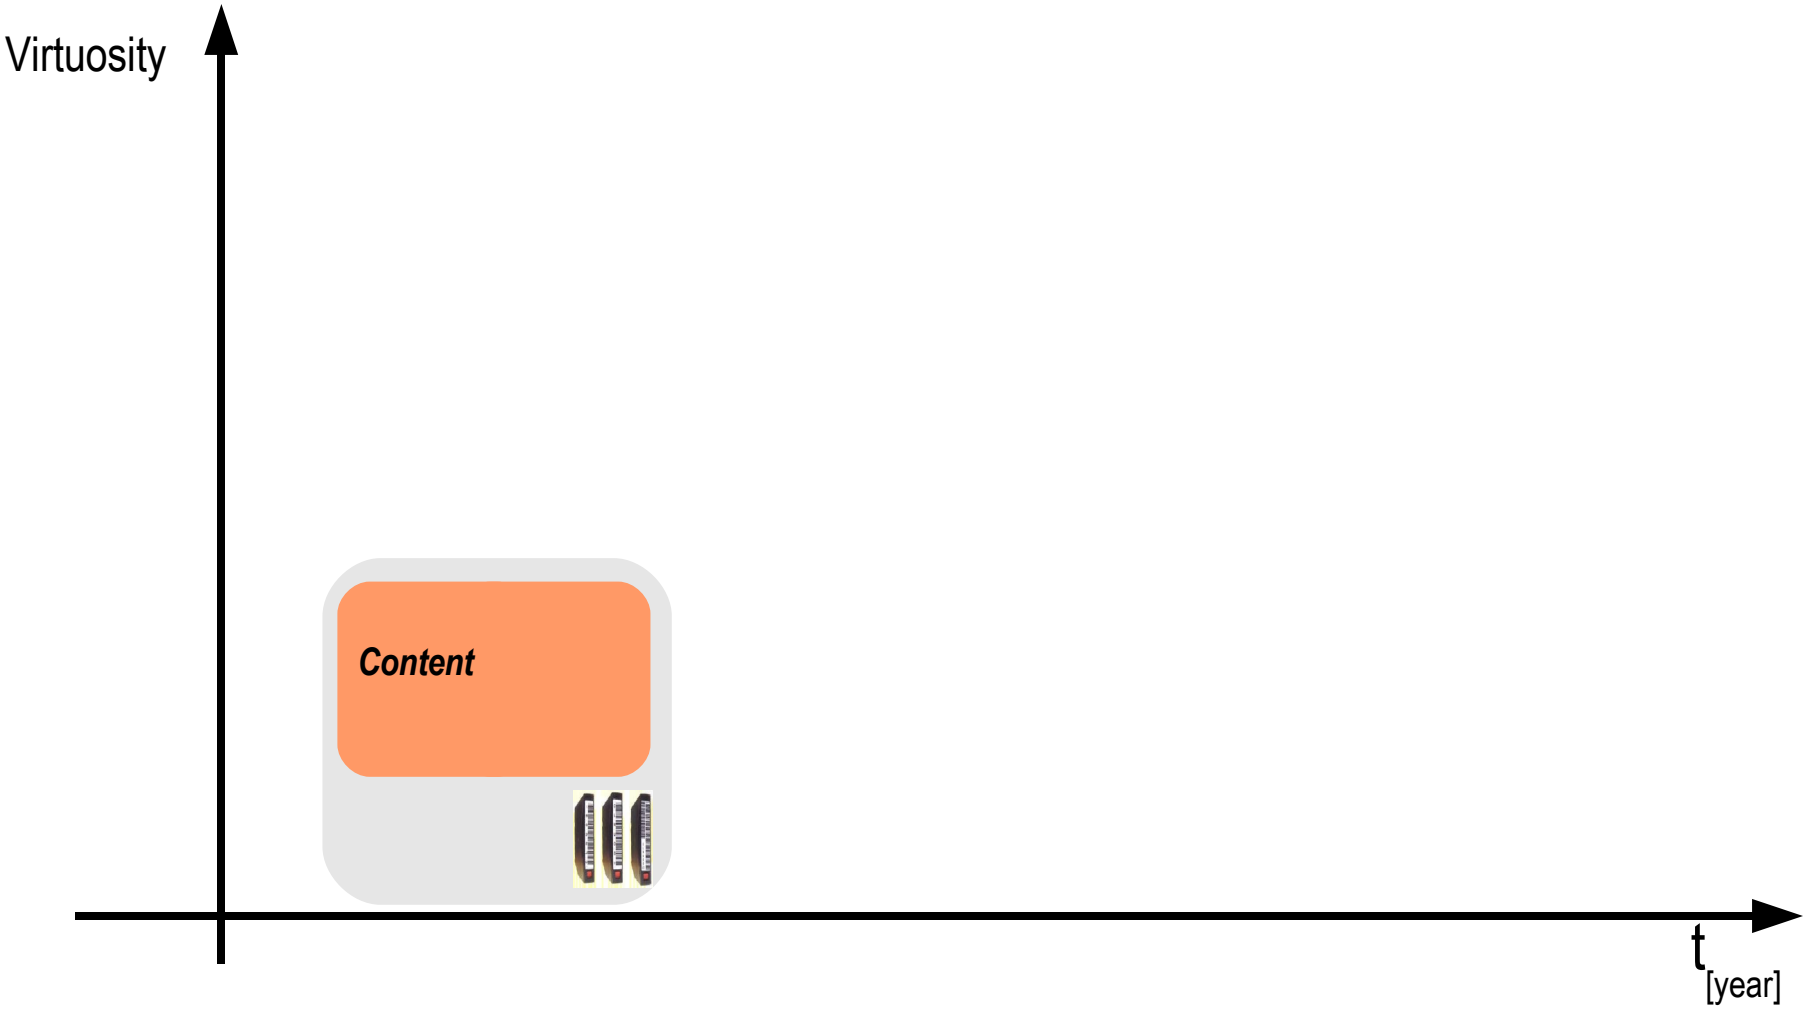

Principle of Migration

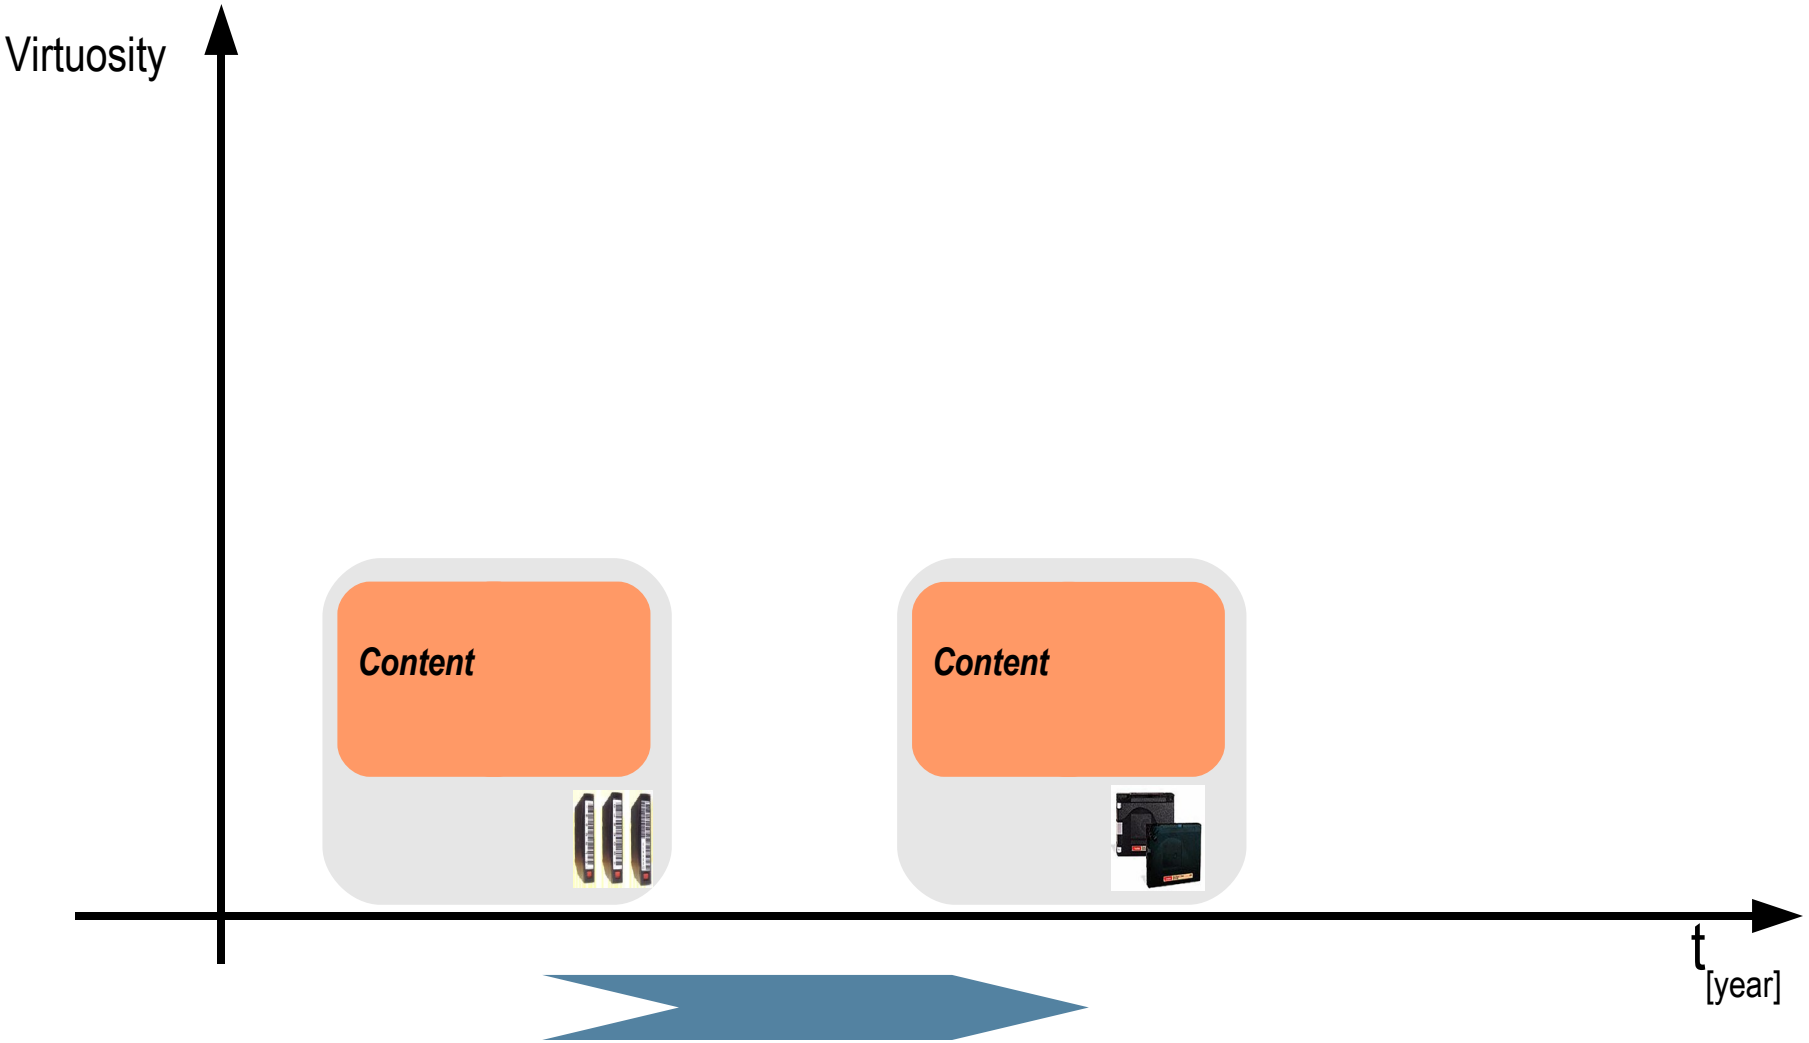

Principle of Migration

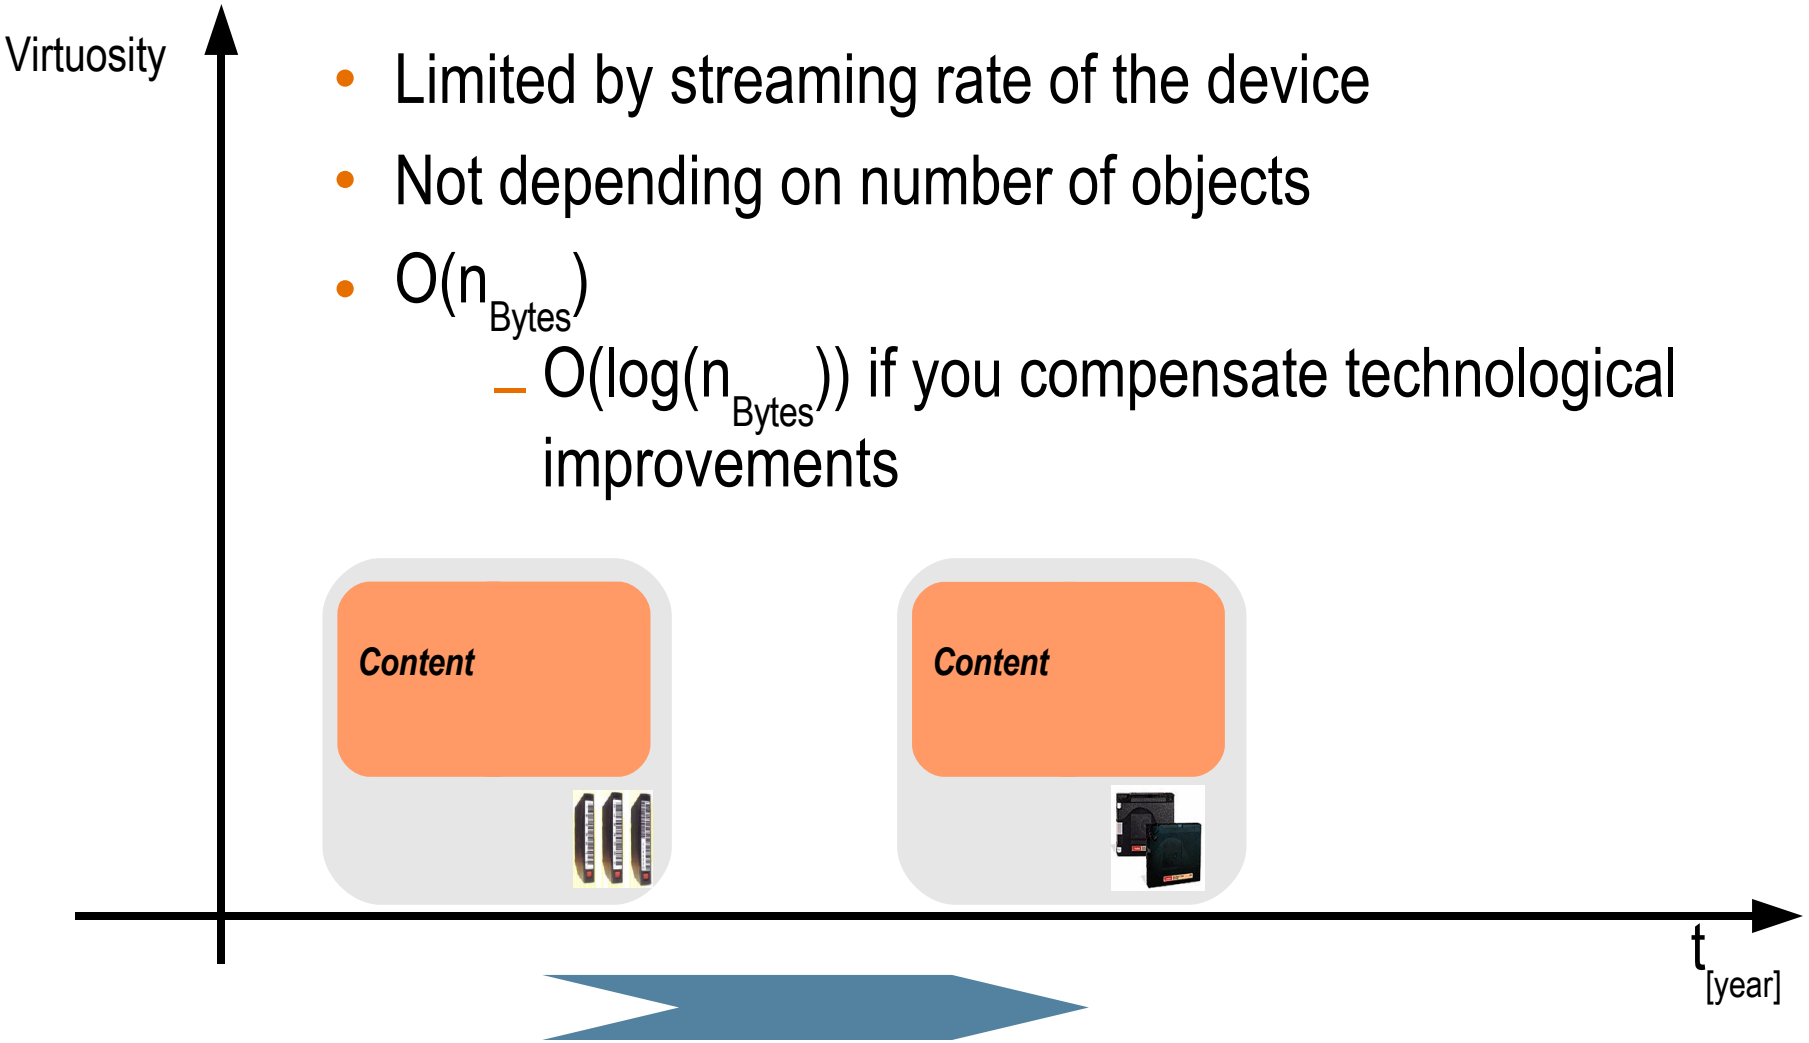

- Limited by streaming rate of the device
- Not depending on number of objects
- $O(n_{Bytes})$
  - $O(\log(n_{Bytes}))$ if you compensate technological improvements

Principle of Migration

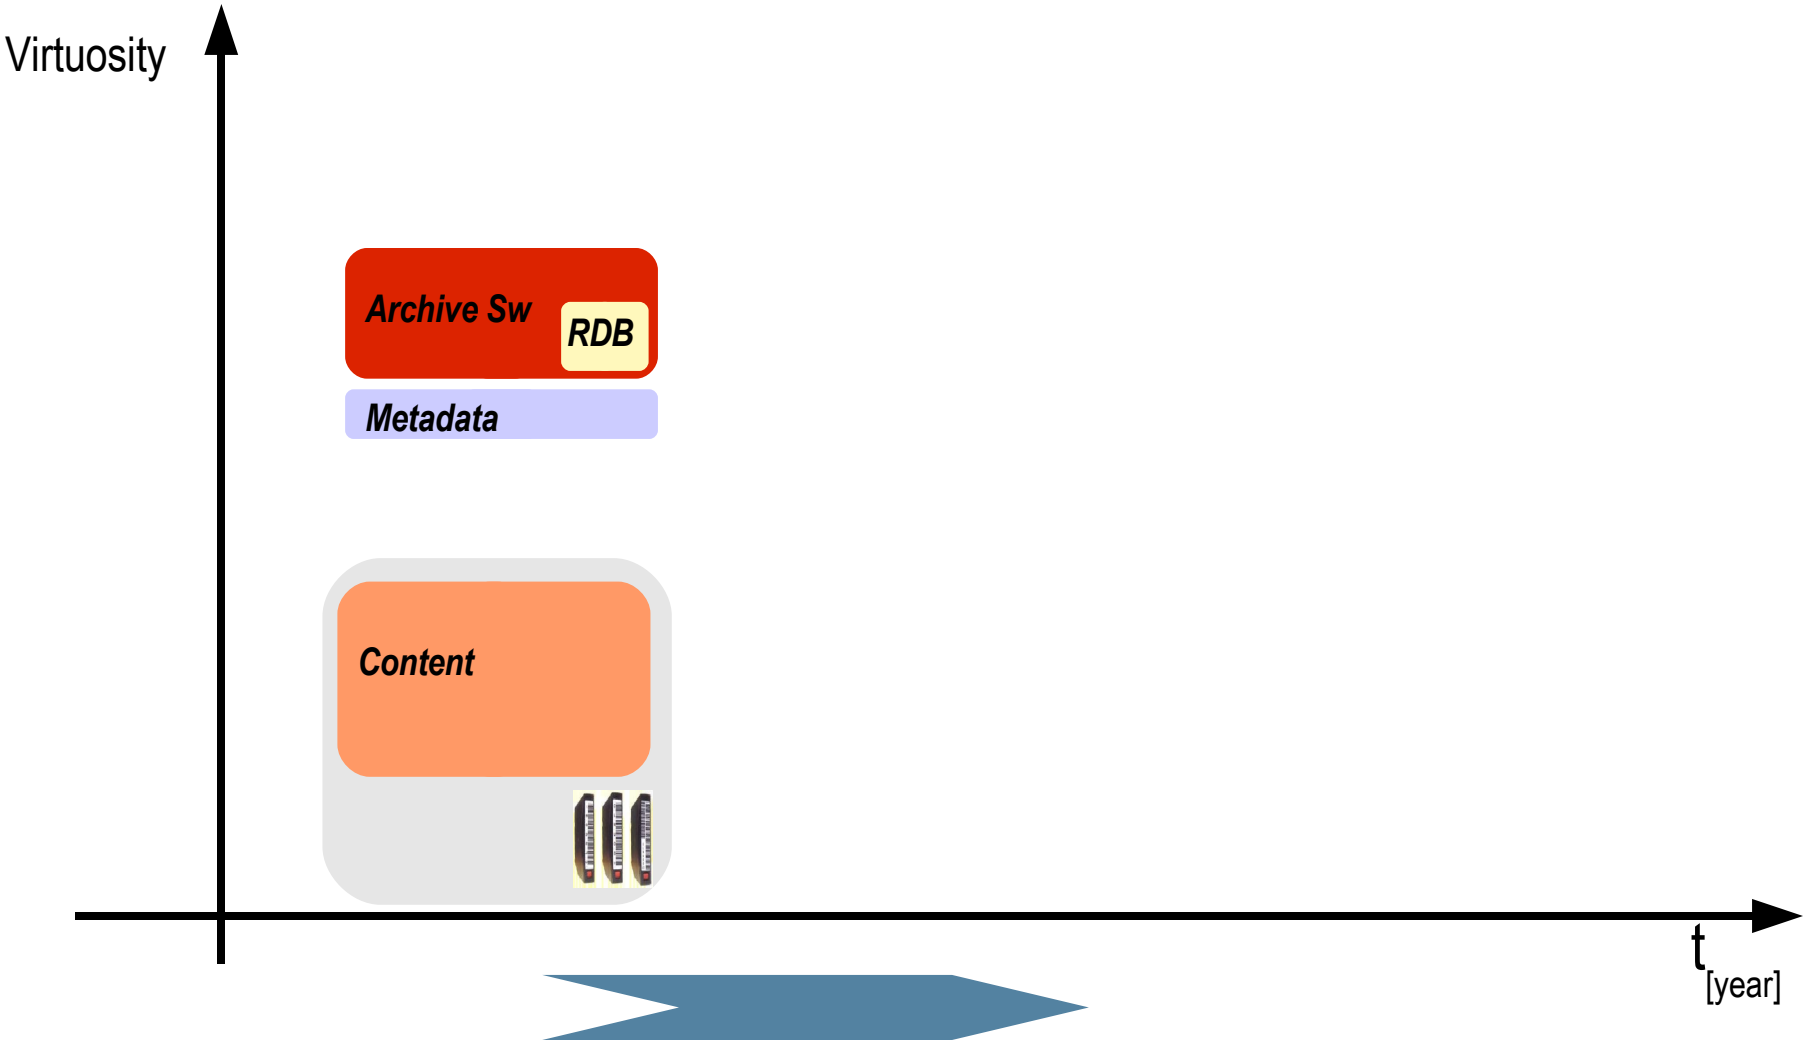

Principle of Migration

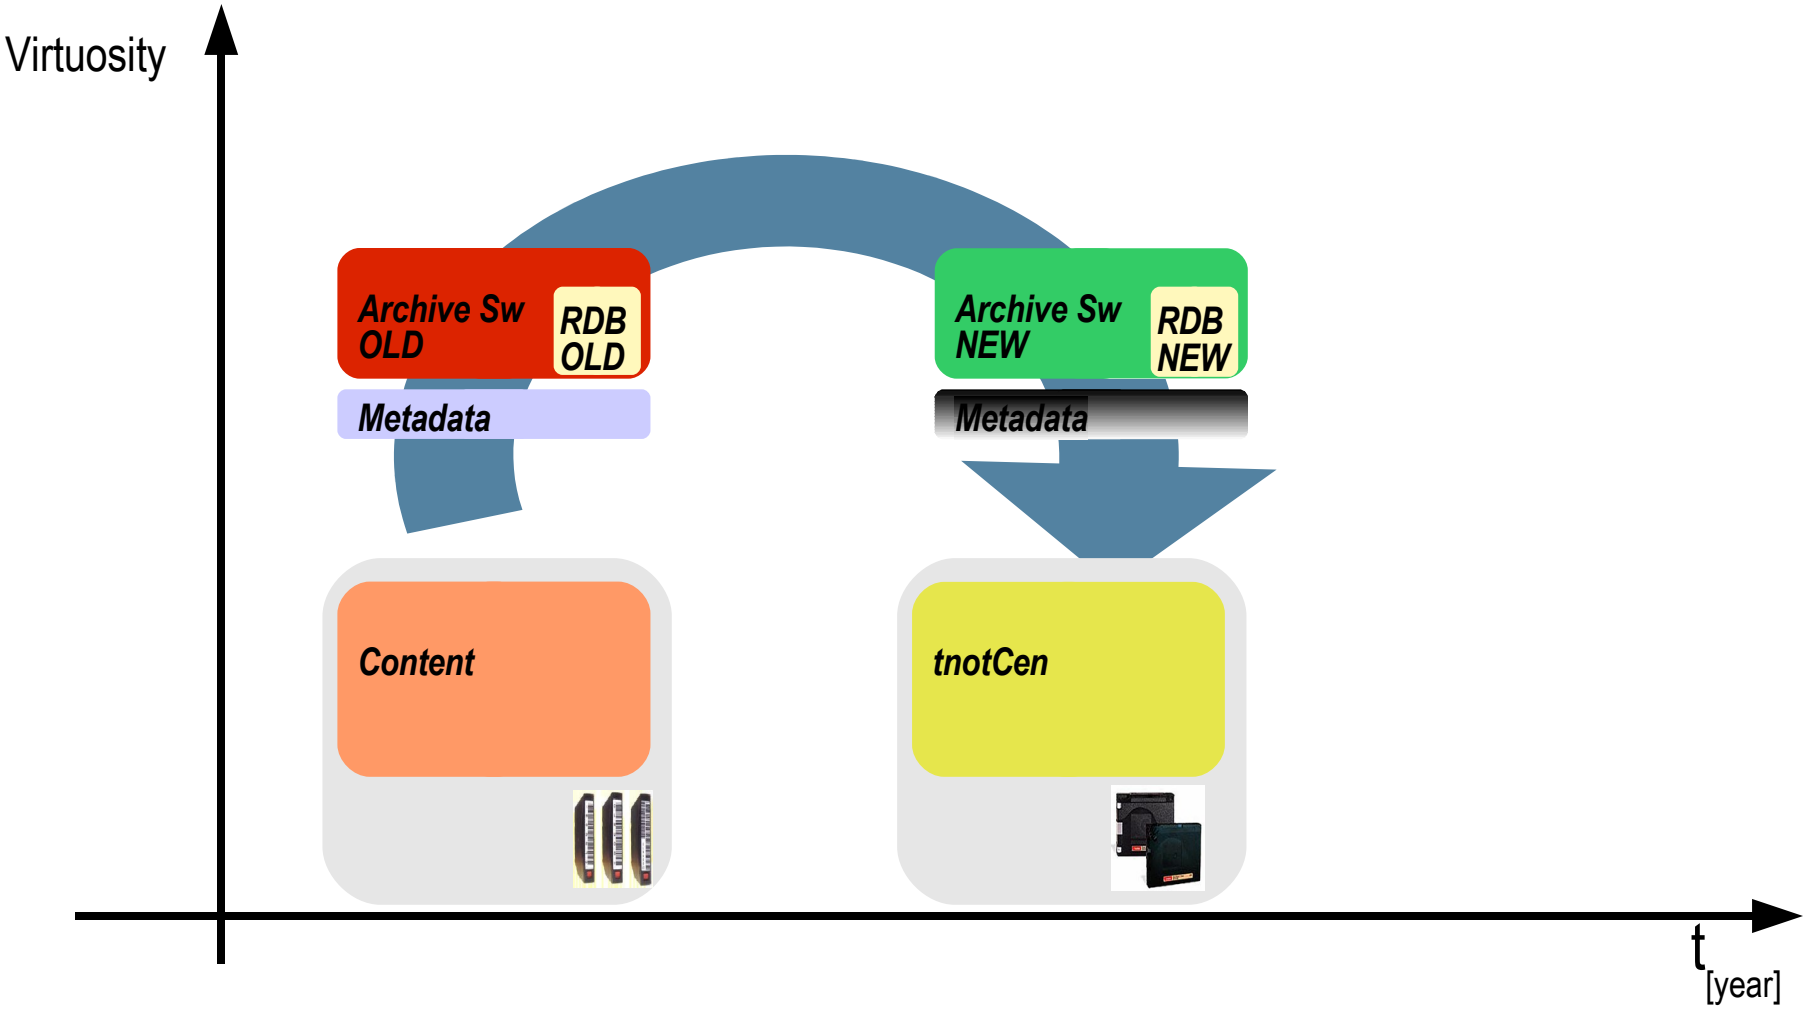

Principle of Migration

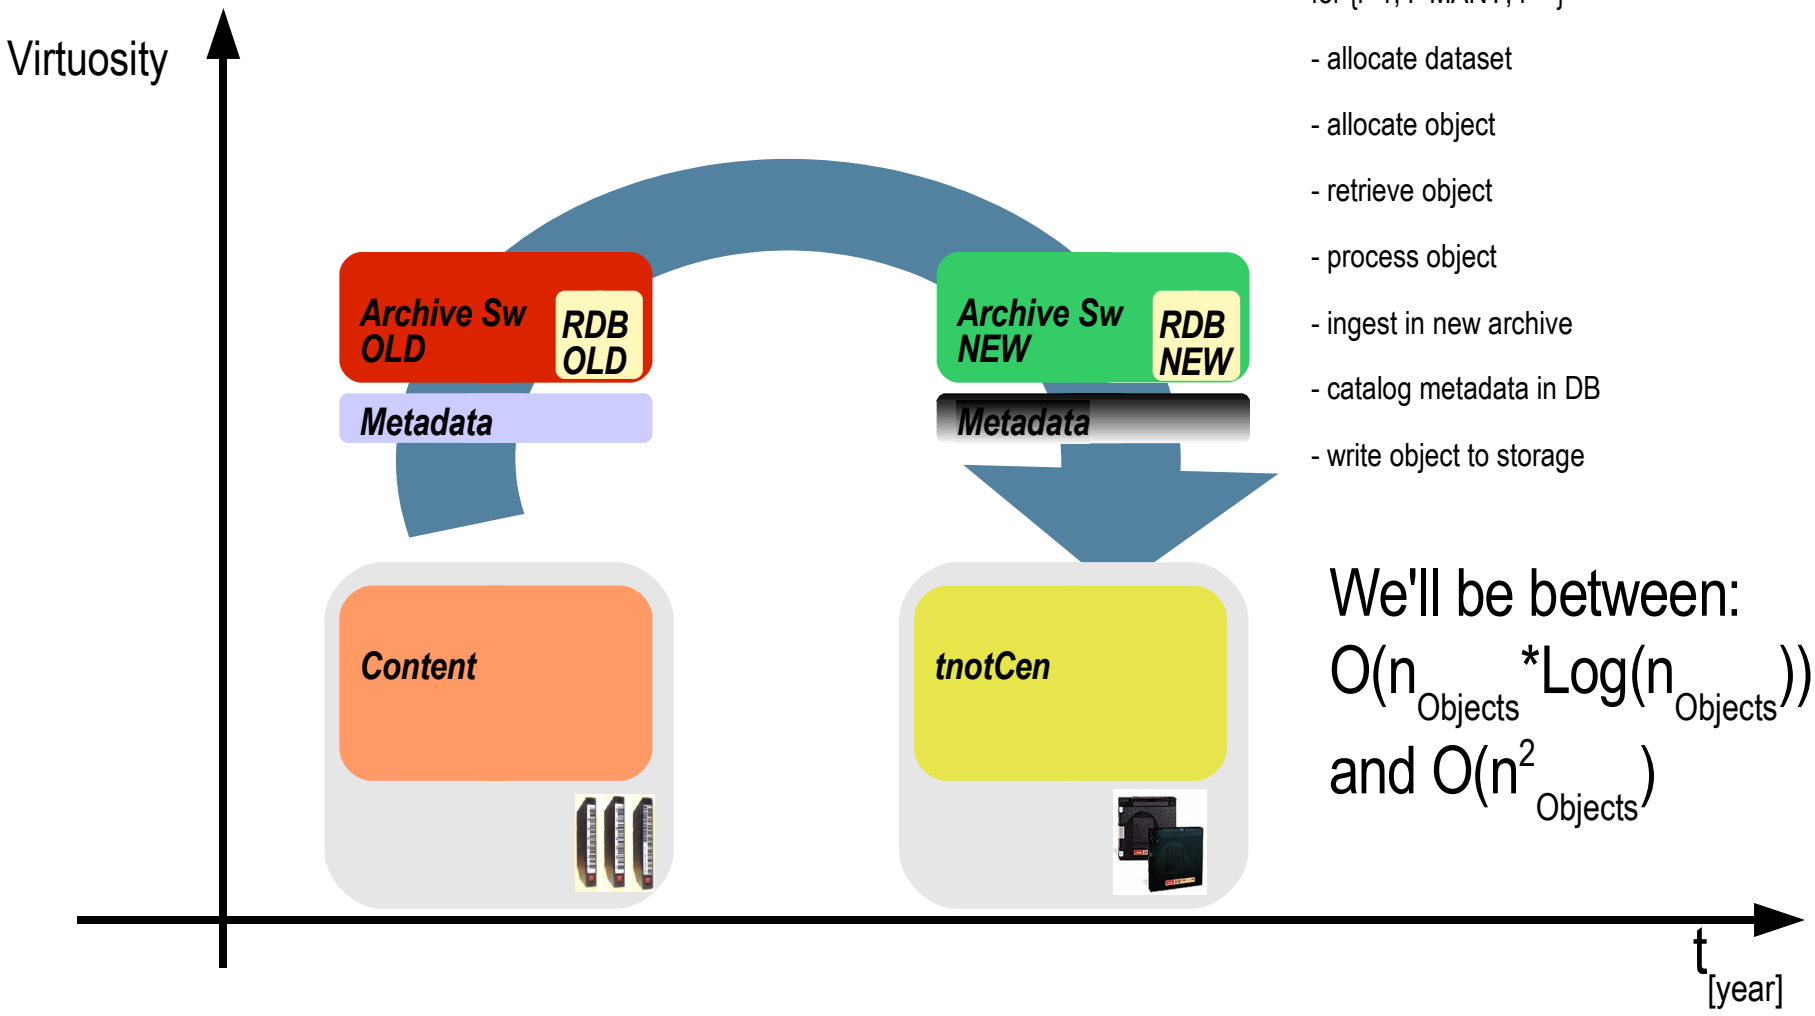

for {i=1; i=MANY; i++}

- allocate dataset
- allocate object
- retrieve object
- process object
- ingest in new archive
- catalog metadata in DB
- write object to storage

We'll be between:

$O(n_{Objects} * Log(n_{Objects}))$

and $O(n^2_{Objects})$

Just an example for AD 2039

- I started with 100million objects. Today,
  - One avarage object is 1MB large (-> 100 TB total)
  - One drive at 100MB/s, one requester at 100 /s
  - One drive can walk my content in 11.4 days
  - Can I retrieve all objects at the same speed?

Just an example for AD 2039

- With a **20%** annual growth, AD 2039
  - I have ingested 28.5 billion objects
  - One avarage object is 10MB large (-> 285 PB total)
  - Current storage devices stream at 6TB/s
  - Migrating all media takes half a day (with one drive)
  - How many objects can I access at random?
    - If I follow $O(n_{Objects} * Log(n_{Objects}))$, it's 13,000 days
    - $O(n^2_{Objects})$ will get me to 925,965 days = 2,535 years
  - Assuming 100% same level of technical progress (streamimg was factor 60,000 faster)
    - Between 0.2 days
    - And 15.4 days

Another example for AD 2039

- With a **50%** annual growth, AD 2039
  - I have ingested 19,175 billion objects
  - One avarage object is 10MB large (-> 191,750 PB total)
  - Current storage devices stream at 6TB/s
  - Migrating all media takes 60 days (with one drive)
  - How many objects can I access at random?
    - If I follow $O(n_{Objects} * Log(n_{Objects}))$, it's 35,900 years
    - $O(n^2_{Objects})$ will get me to 419 billion years
  - Assuming 100% same level of technical progress (streamimg was factor 60,000 faster)
    - Between 0.6 years
    - And 7 million years

Re-Use in the Far Future?

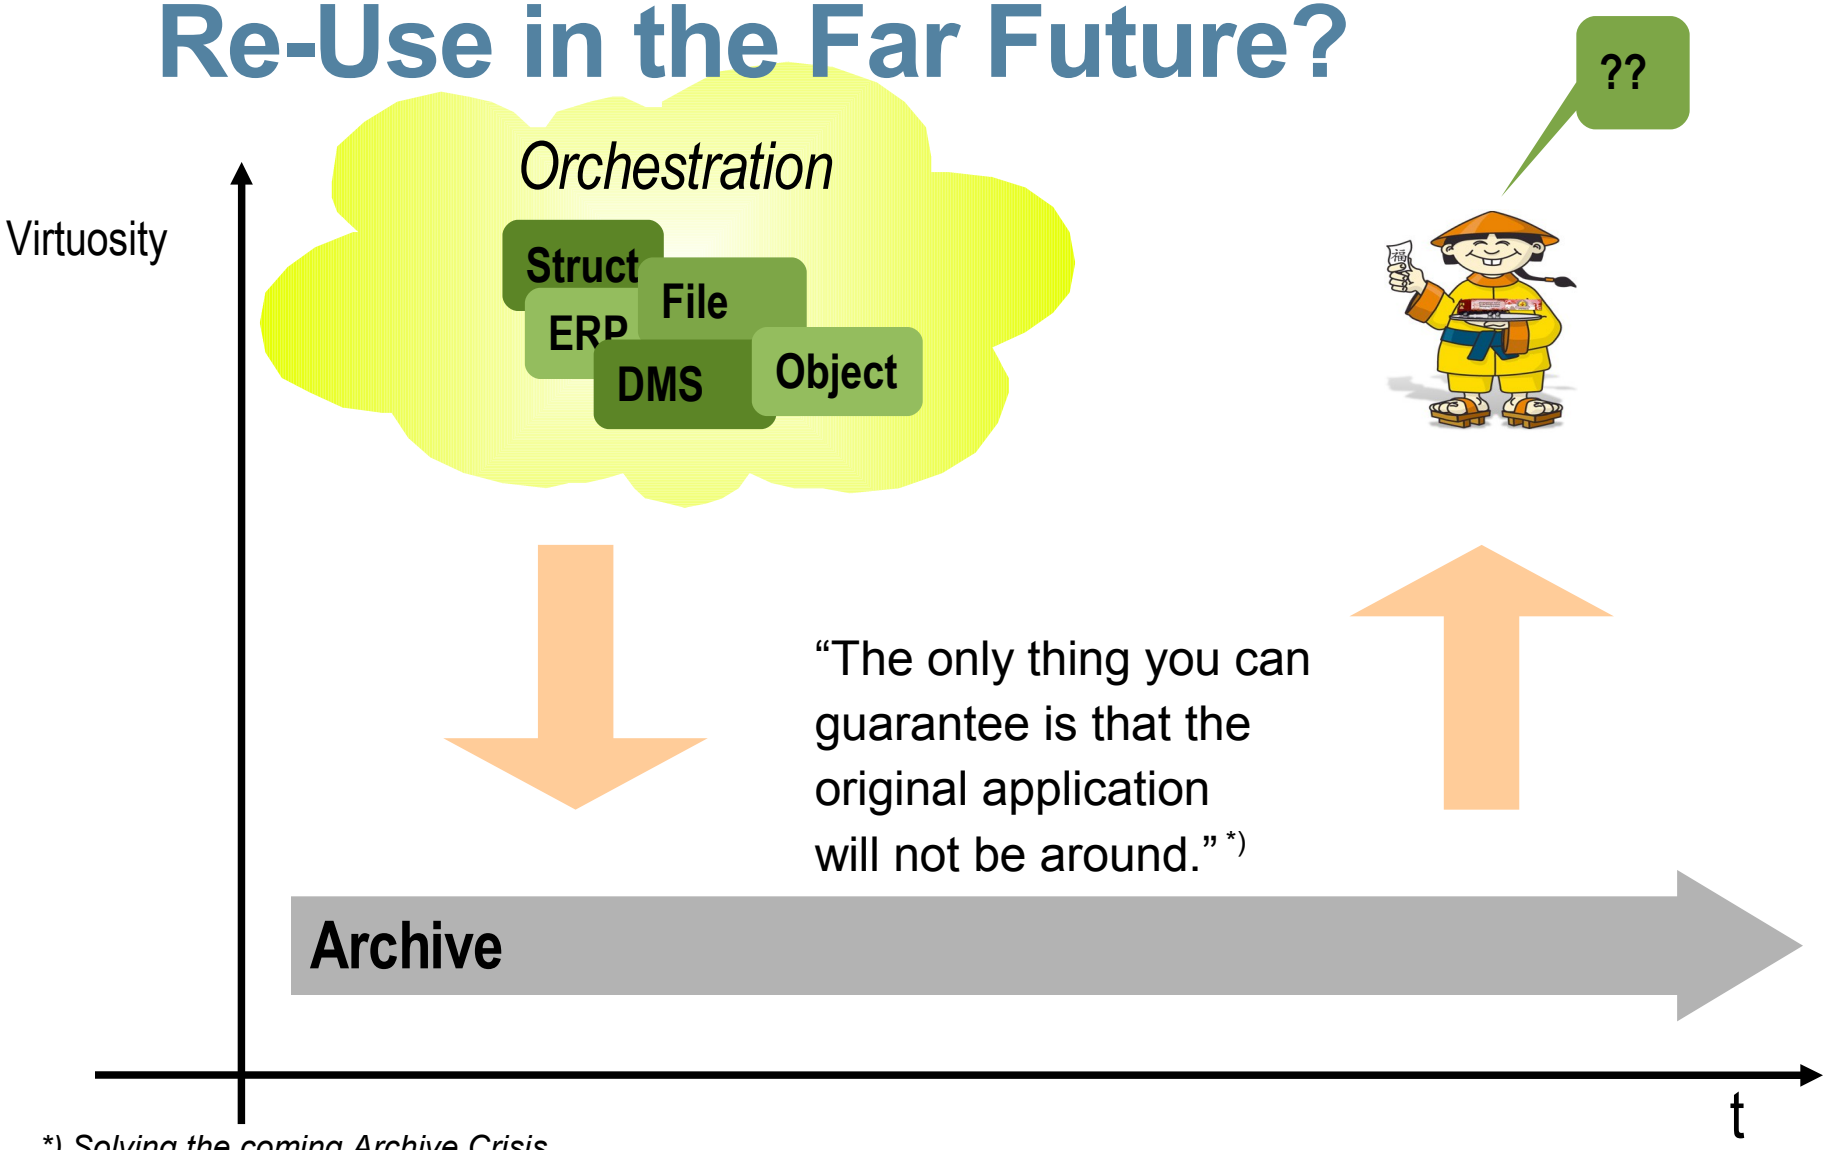

*) Solving the coming Archive Crisis
http://www.snia.org/forums/dmf/programs/ltacsi/100_year/

Yet Another Migration Currency

- 1 Pimb (= 1 Pain-in-my-back)
- Equals 1 inherited archive vault
- Is the necessary individual attention and care.
- Is > 90% of the migration effort
- That eats man days!












12 Pimbs

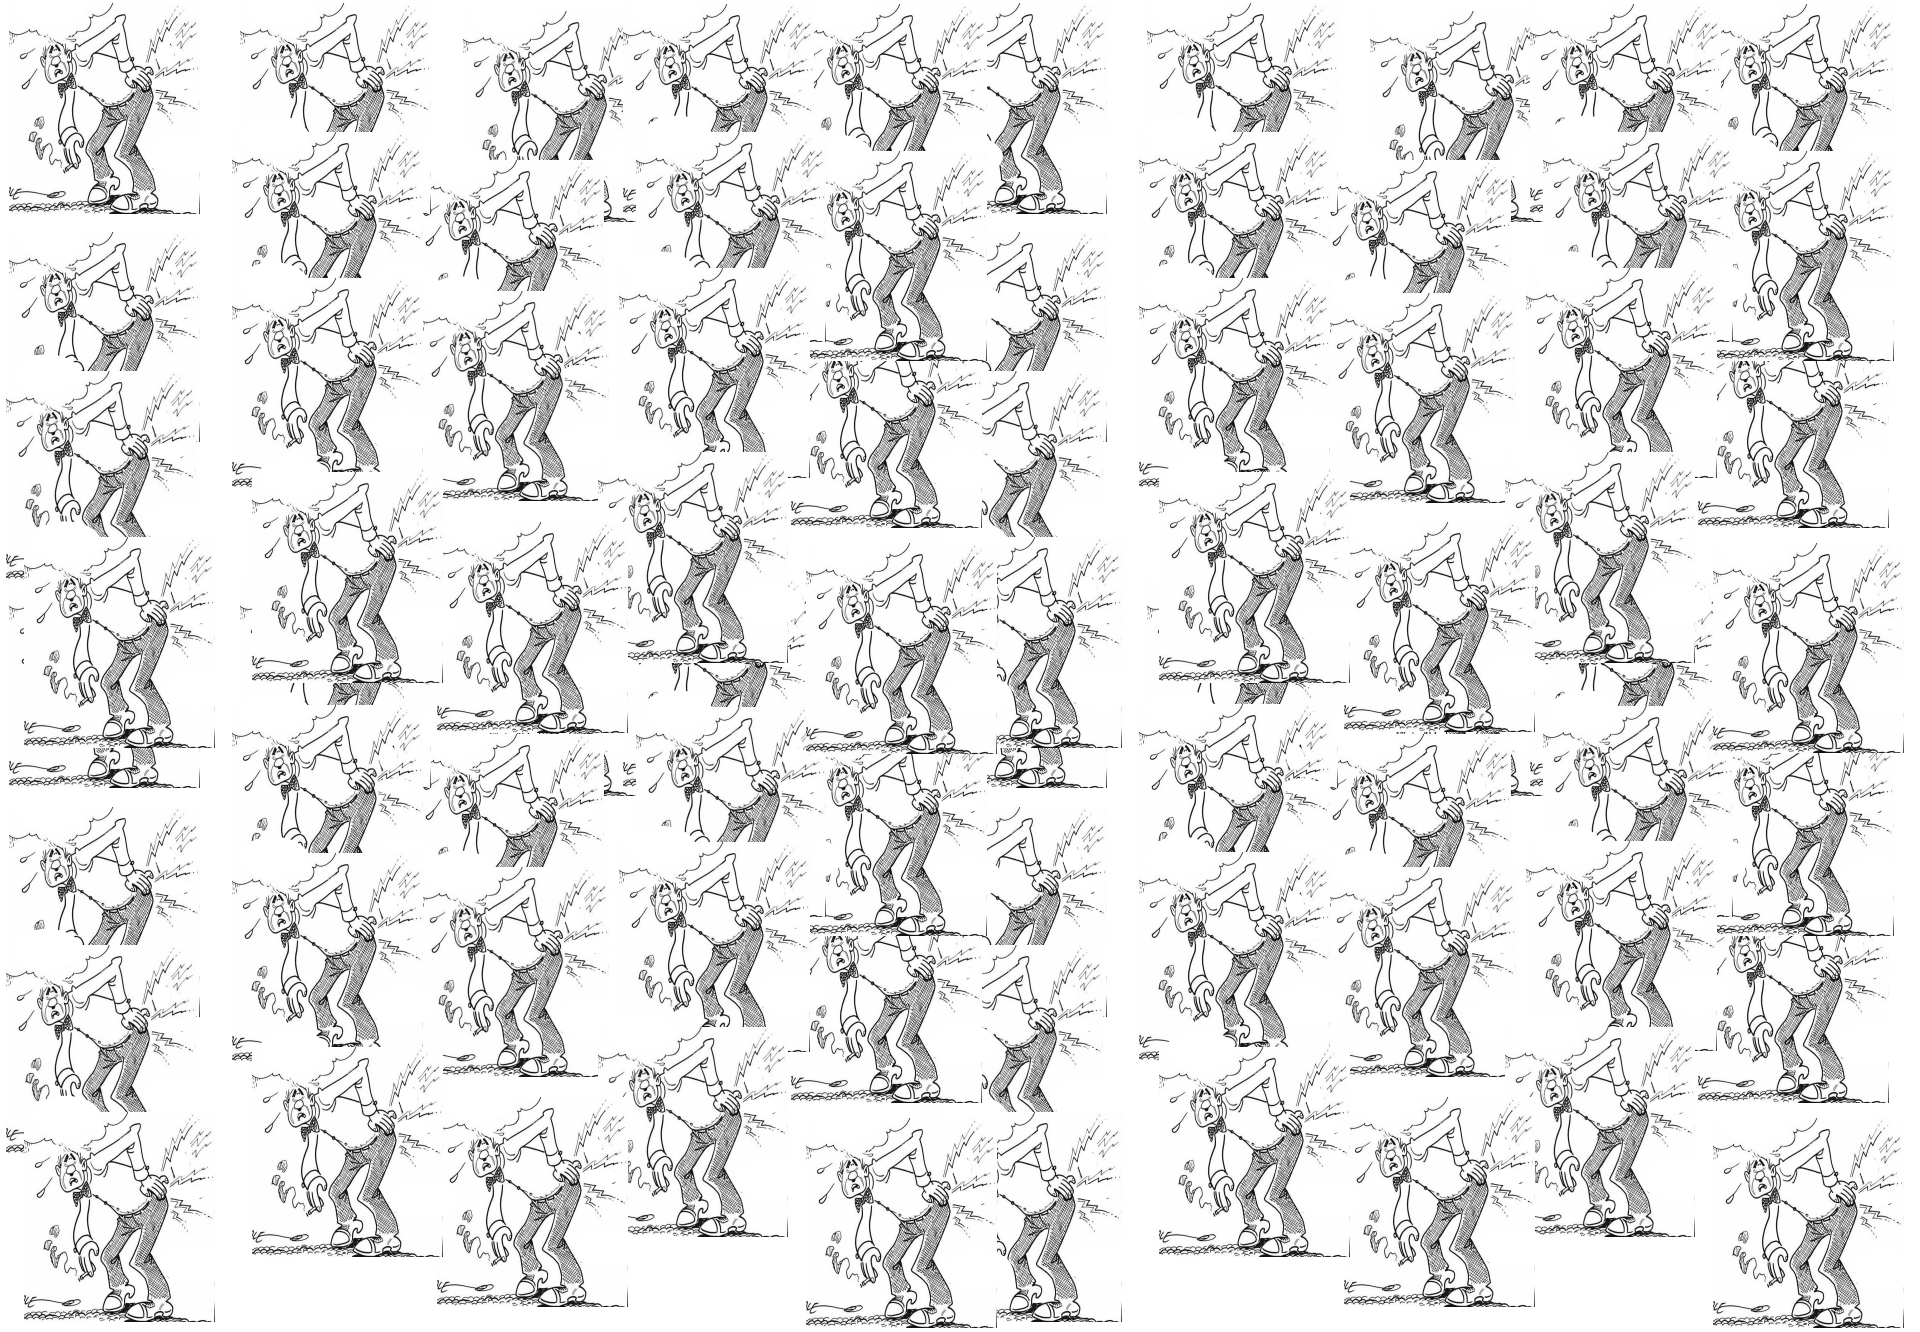

Try a Pimb forecast for AD 2039

- Linear, 1 new sin every 2 years:
  - 27 Pimbs
- Percentual, 5% p.a.
  - 52 Pimbs
- Including a greenness factor:
  - X,000 Pimbs

Take away

- List your components well
- Plan a funeral for every component. There are no hallows.
- Critically look at your orchestration
- Do nice things like content integration (limit Pimbs!)
- I'm not saying it's impossible, but there are some astronomic “oops” factors

Questions?

Thank you!

thorsten.lange@sun.com
http://www.sun.de