In RAM We Trust:
A Modern Approach to Forensic Processing

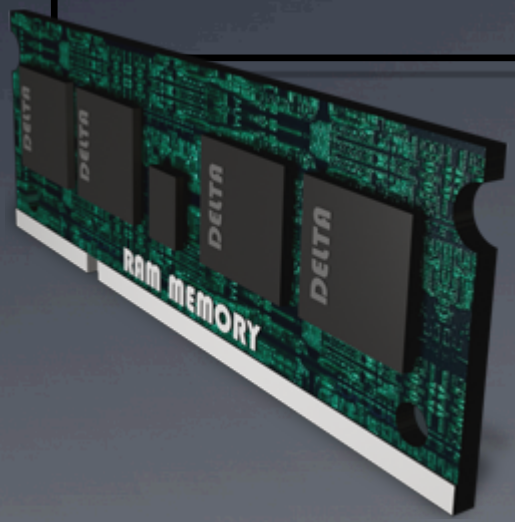

THE UNIVERSITY of NEW ORLEANS

Vassil Roussev

vassil@roussev.net

Review: Primary trends in forensics

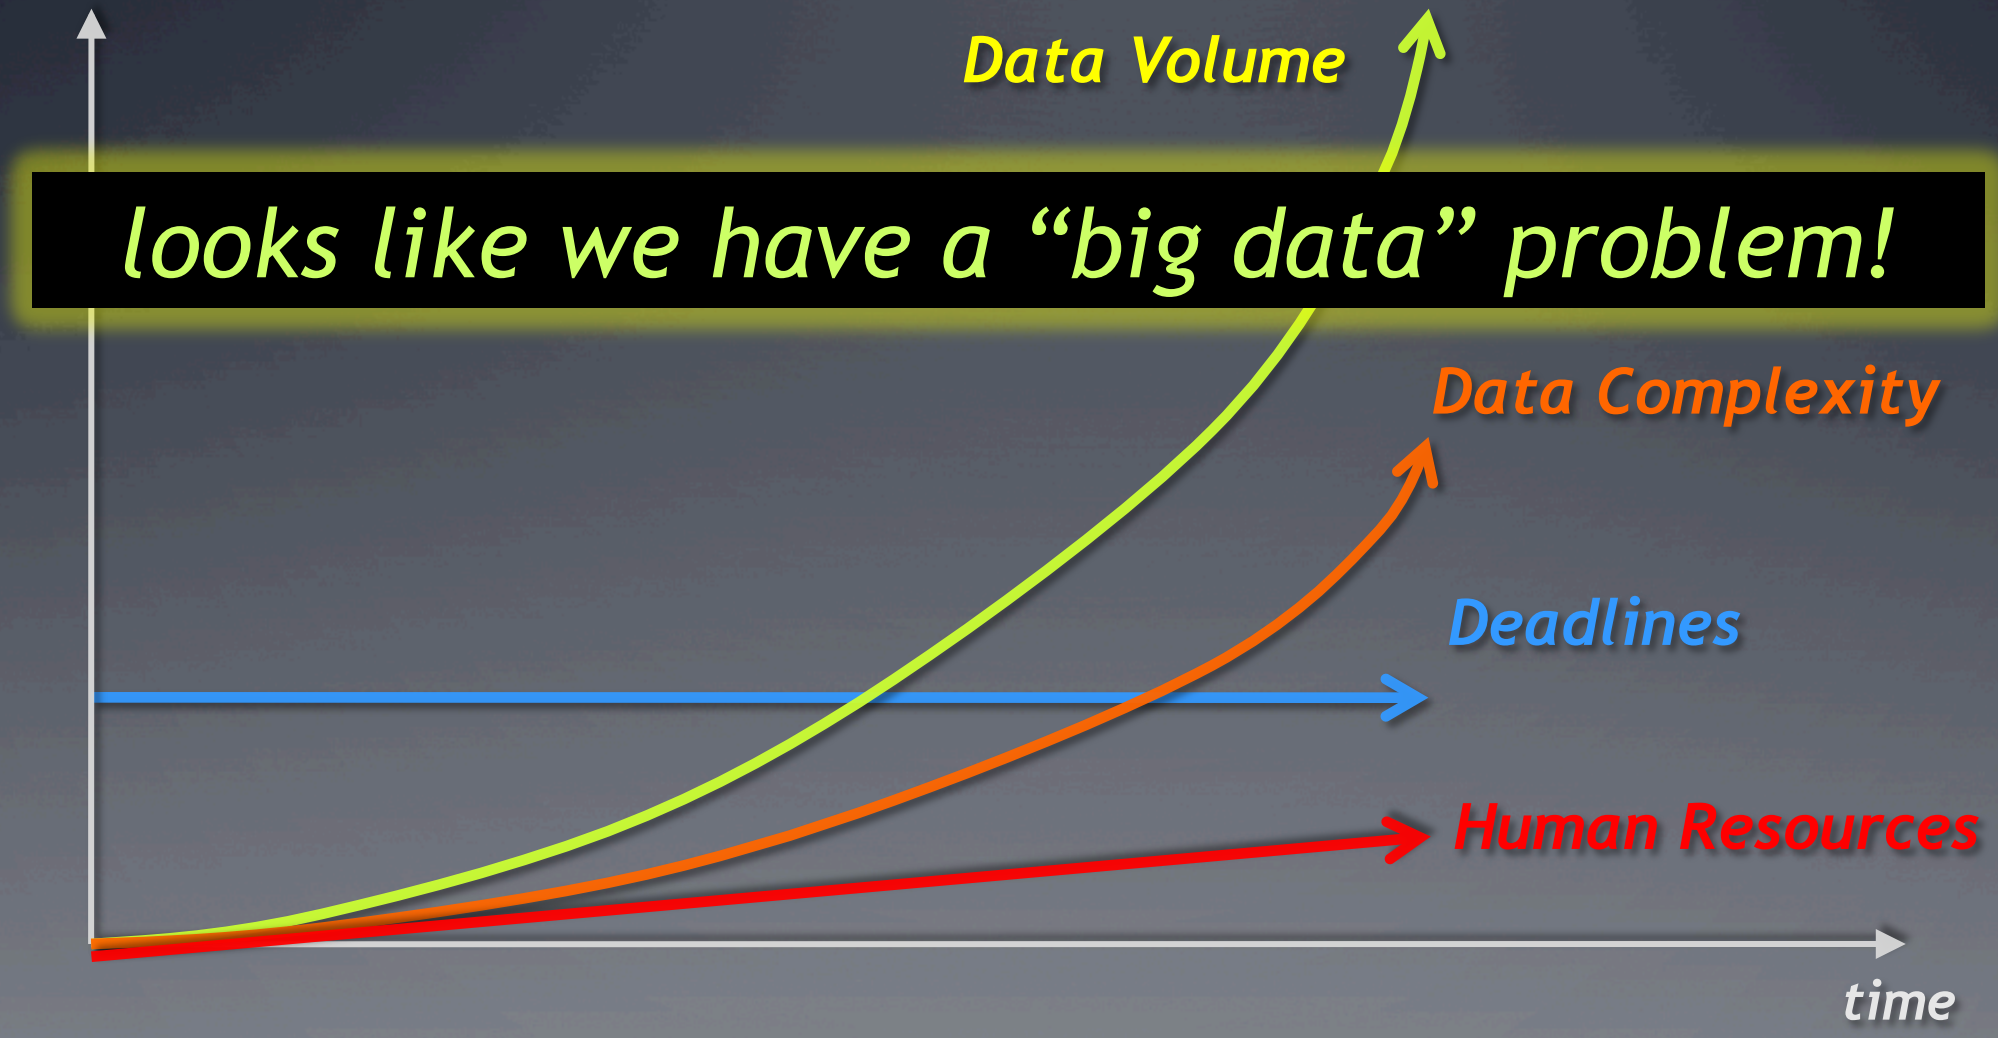

The three "V"s of big data

(by *Michael Stonebraker*)

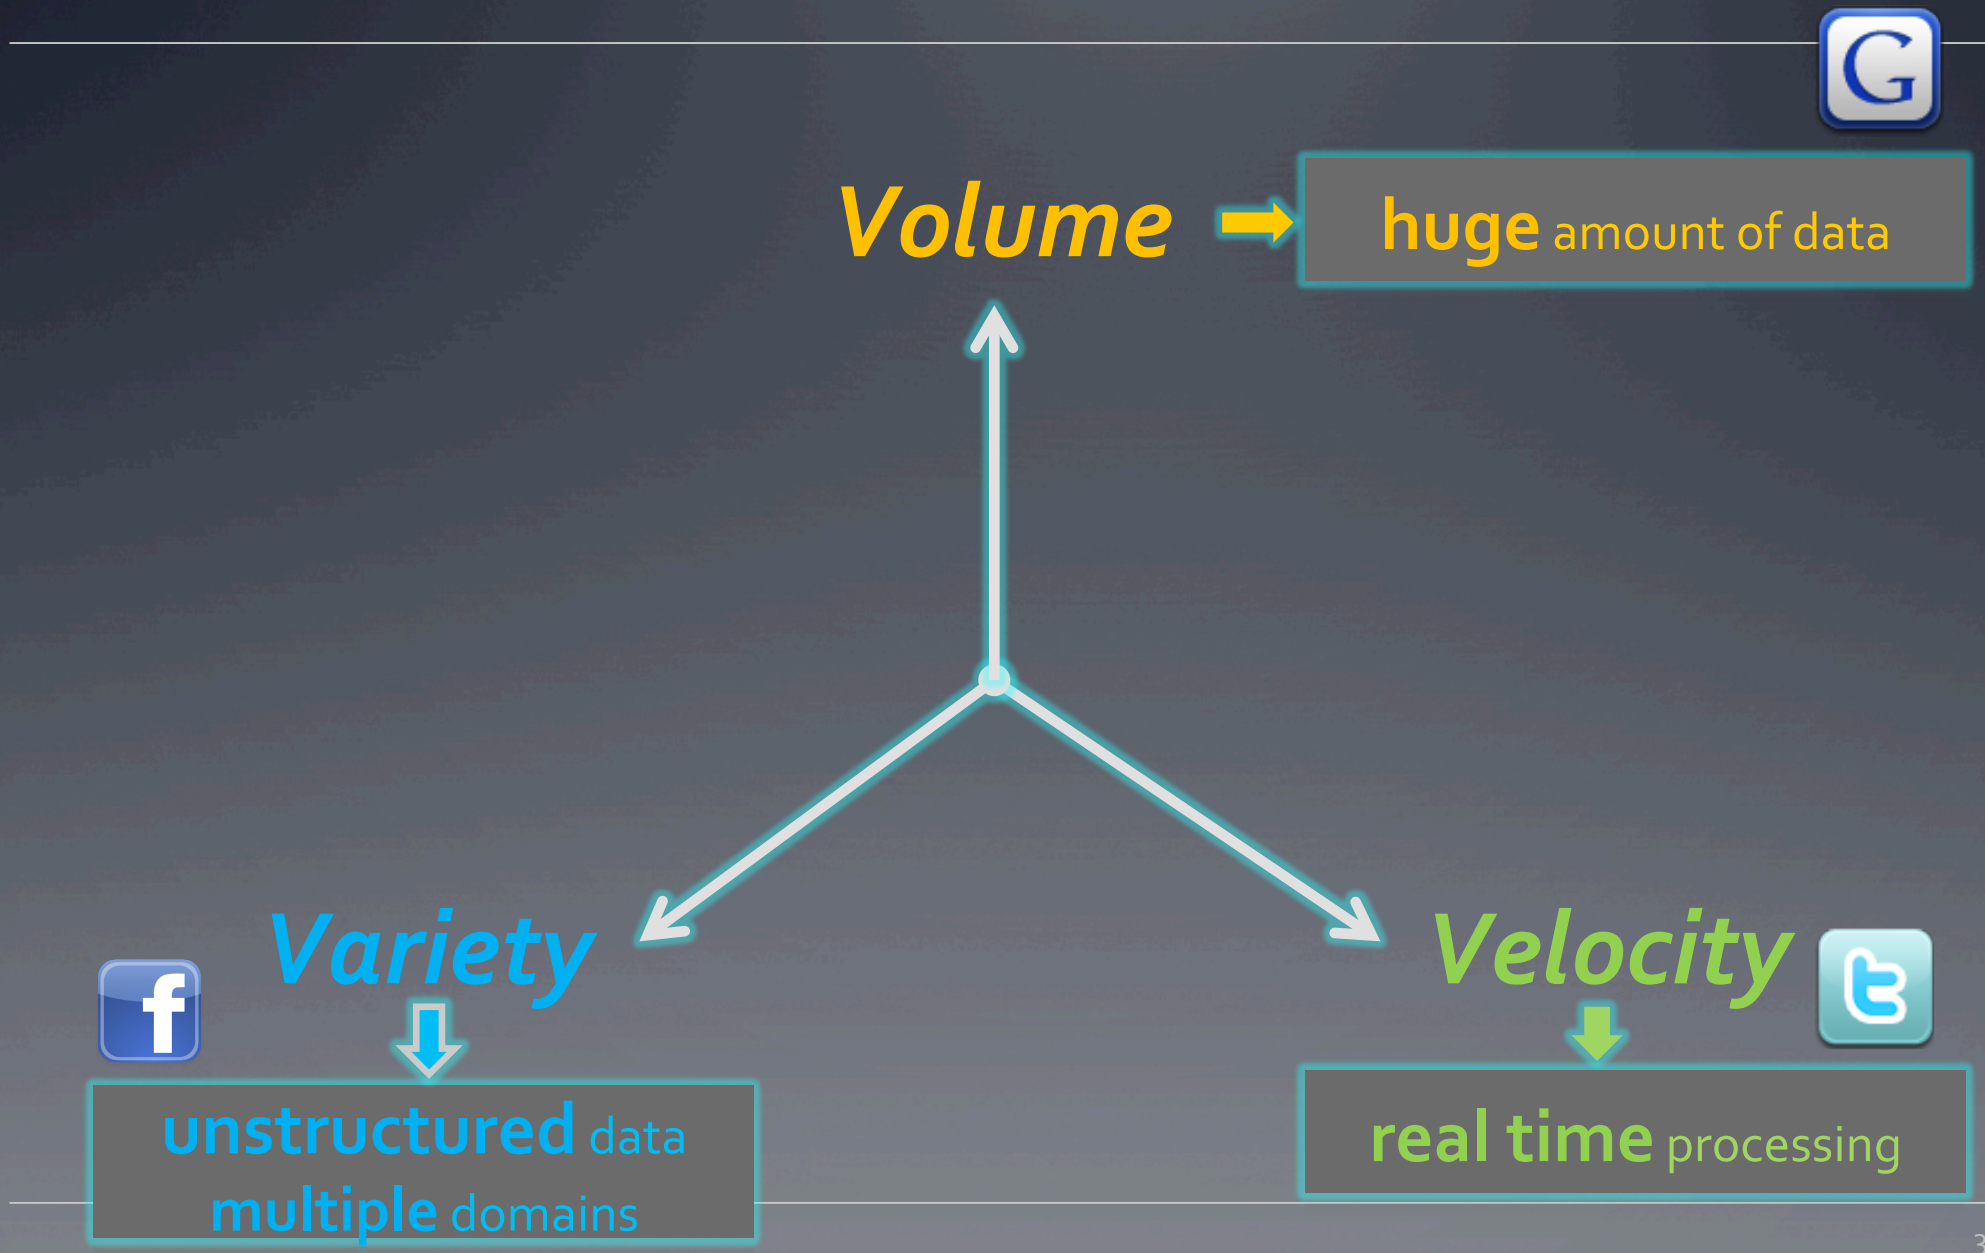

Optimization priorities

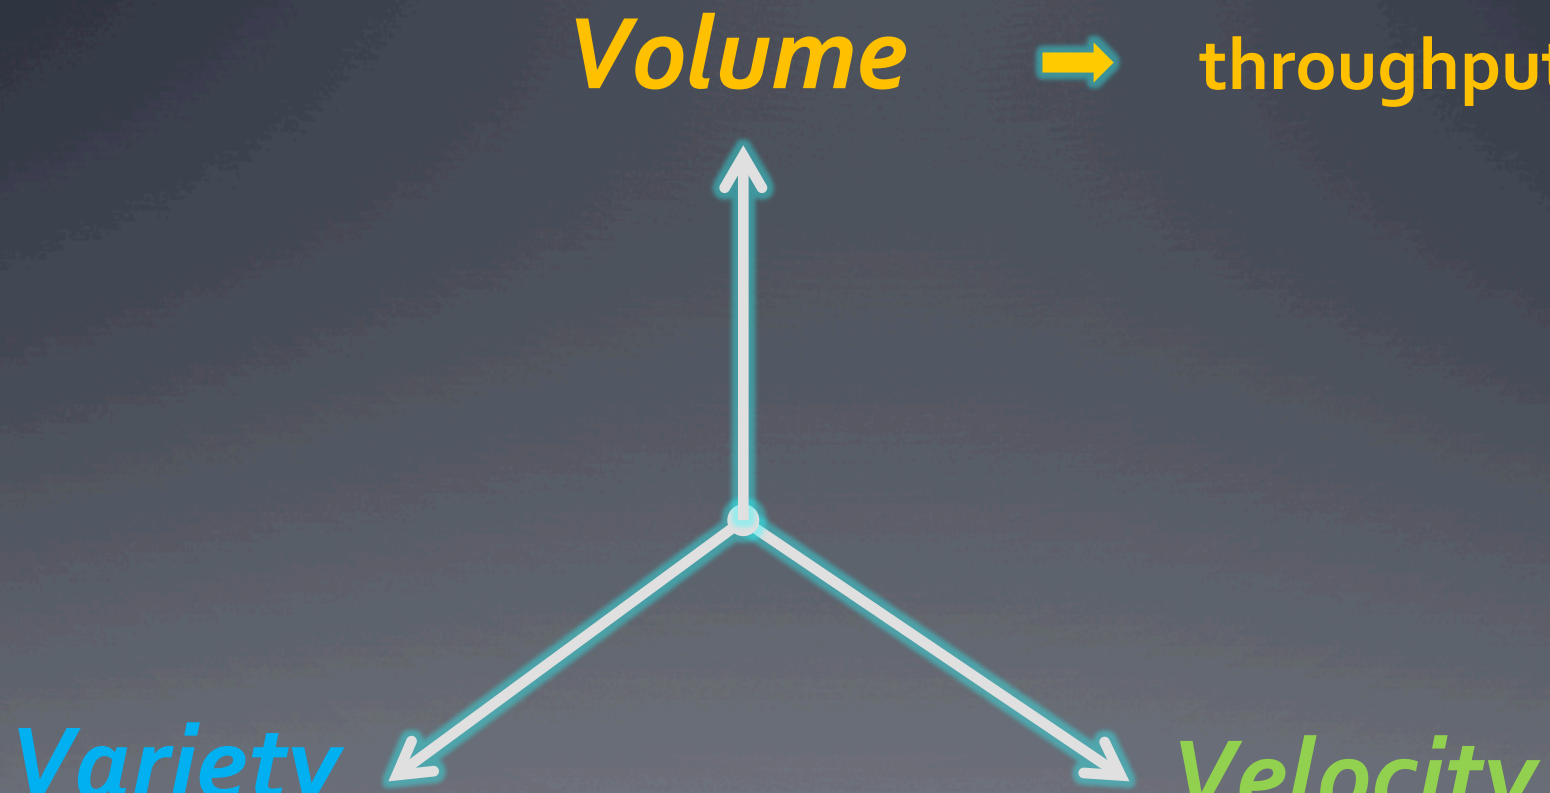

What is forensics' *primary* challenge?

- Volume:
  - Do we have PB of data?

  Not really!

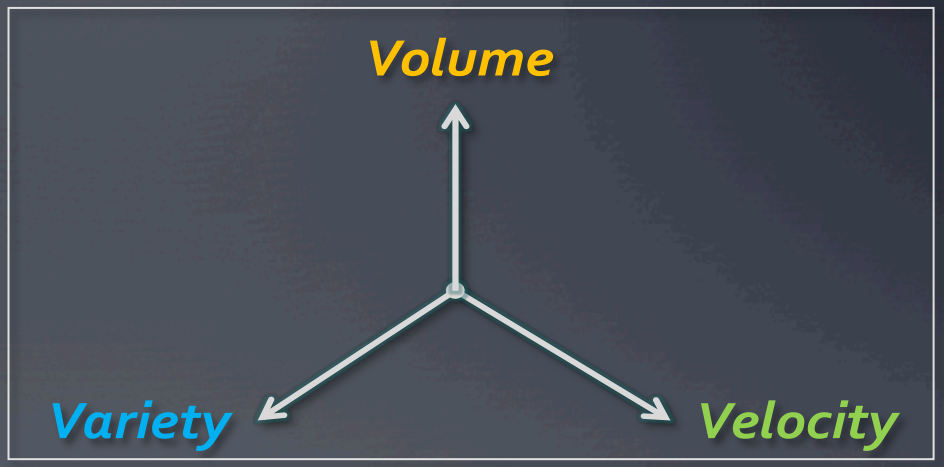

- Variety:
  - How many *types* of things do we need to process?

  Few (but growing)

- Velocity:
  - Do we have deadlines?

  YES!!!

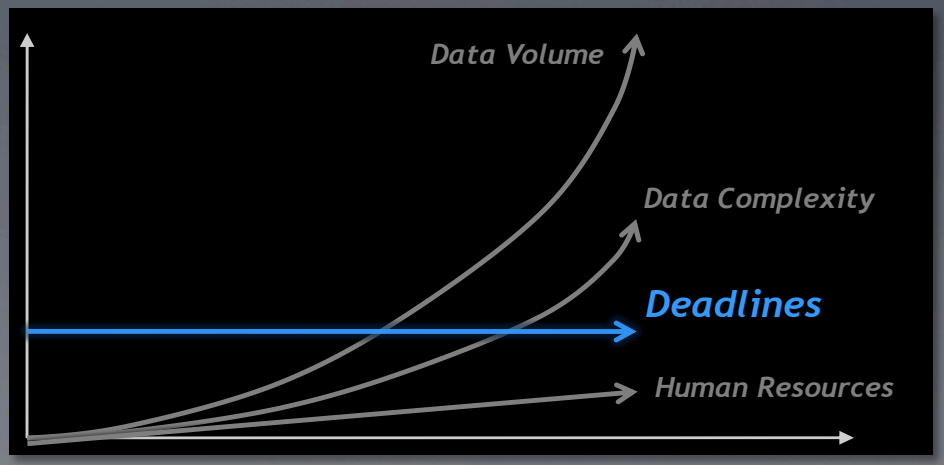

Where are we now?

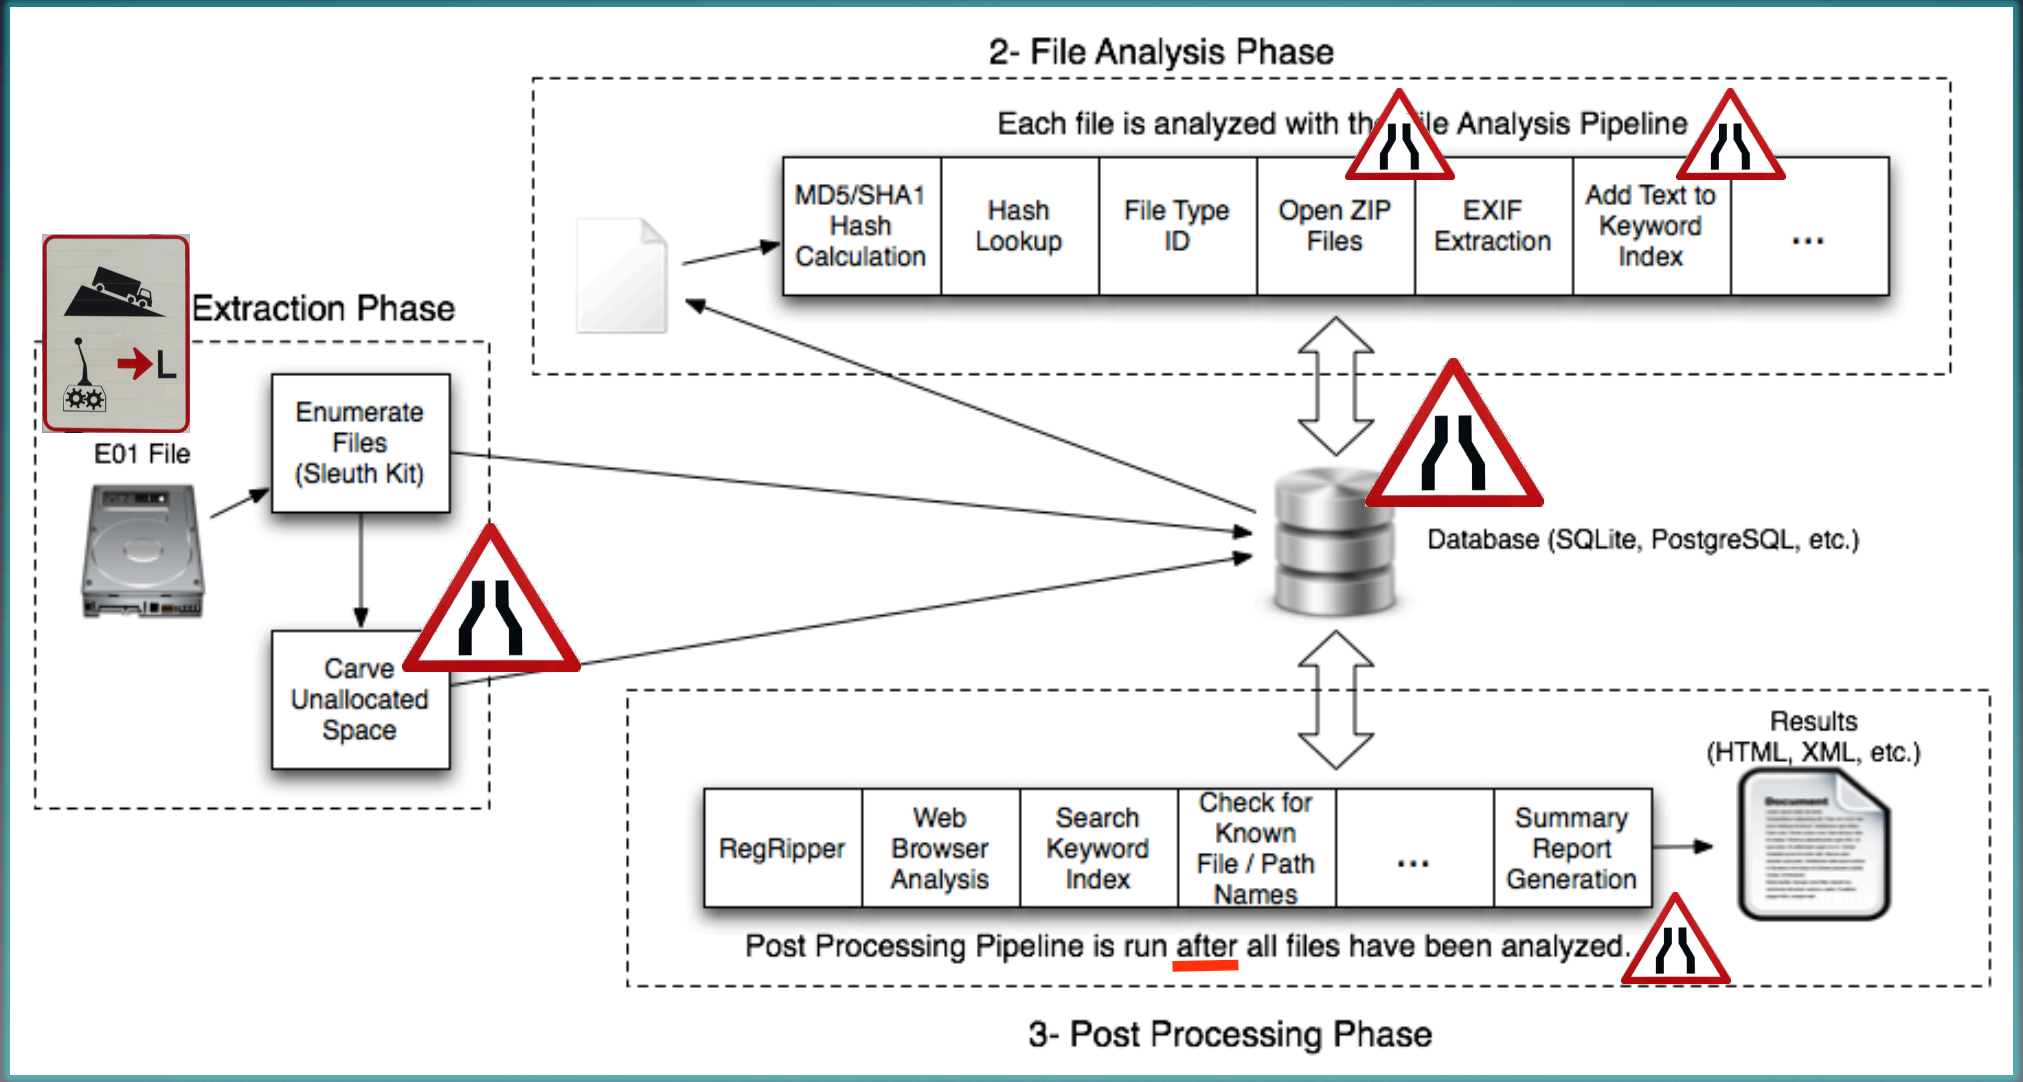

Mapping problems to solutions

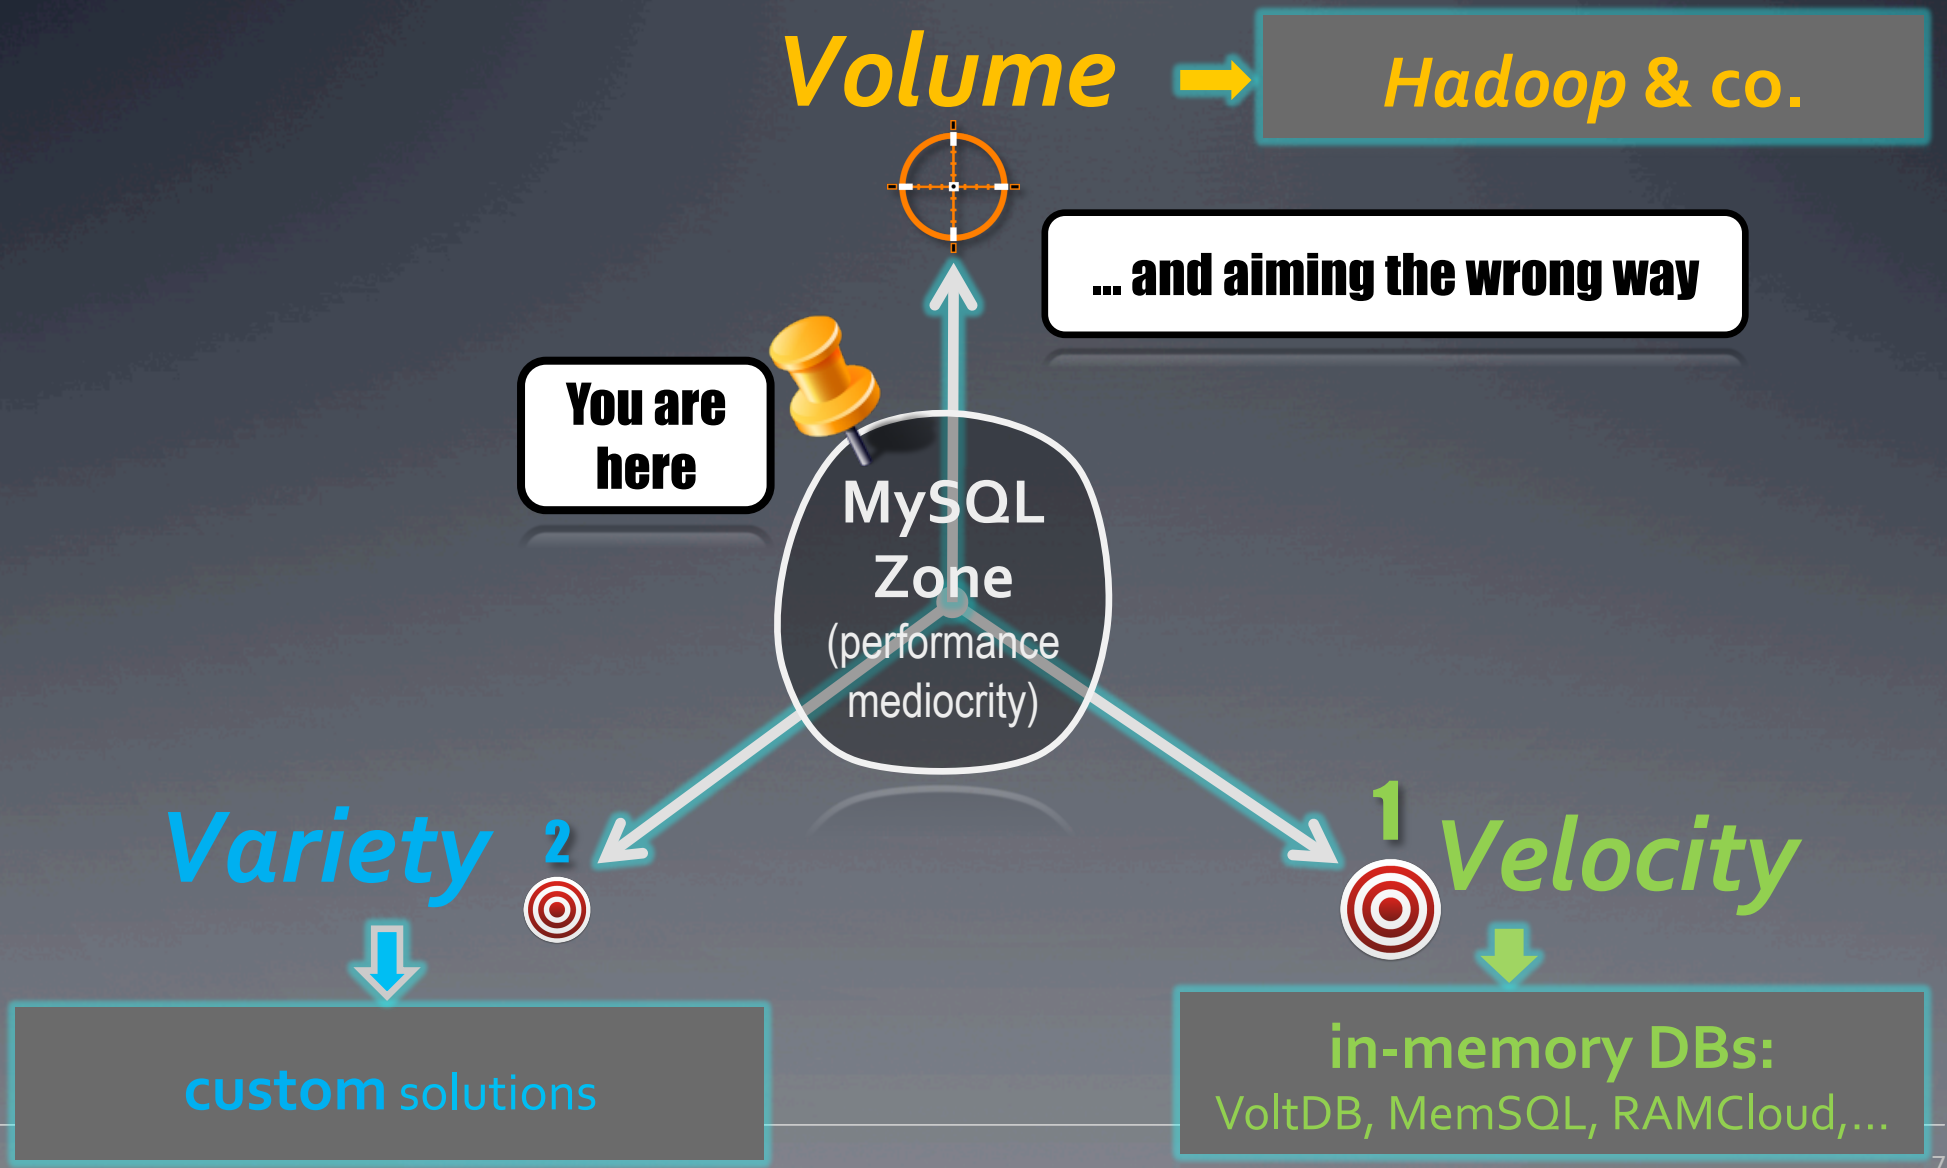

What's wrong with Hadoop?

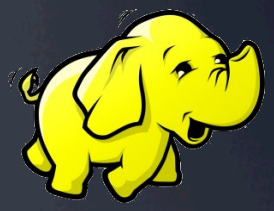

- Nothing
  - … if you have a LOT of data (100TB+)
  - … it's your only choice, really
- If you don't?
  - … you still have to wait

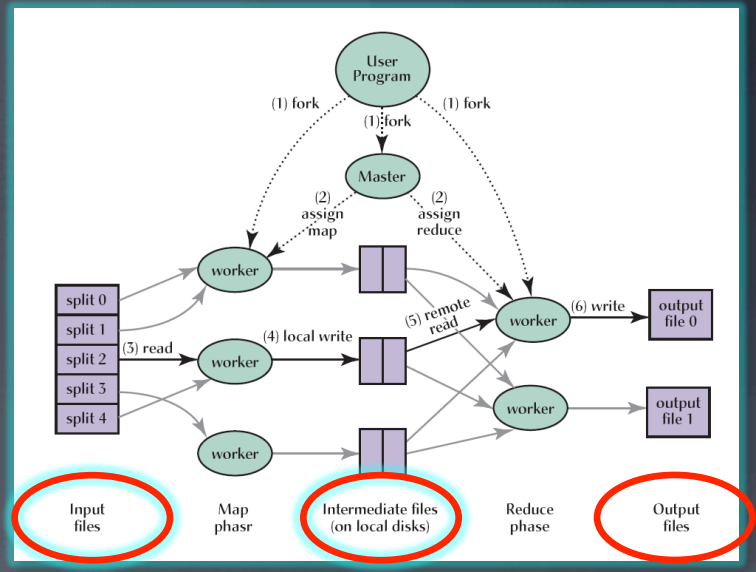

- It is a throughput engine
  - Requires a lot of time to seed initially (HDFS)
  - Suitable for data processing sweeps over *entire* sets
  - Tasks communicate via the file system
  - Not all processing fits the M/R model
  - Will do ***nothing*** to speed up triage and early processing

In other words …

- The failure of current tools to address latency requirements leads to data backlogs.
- This leads to the ***perception*** that we have a volume problem.
- Using a “bigger hammer” designed for volume will do little to address latency.

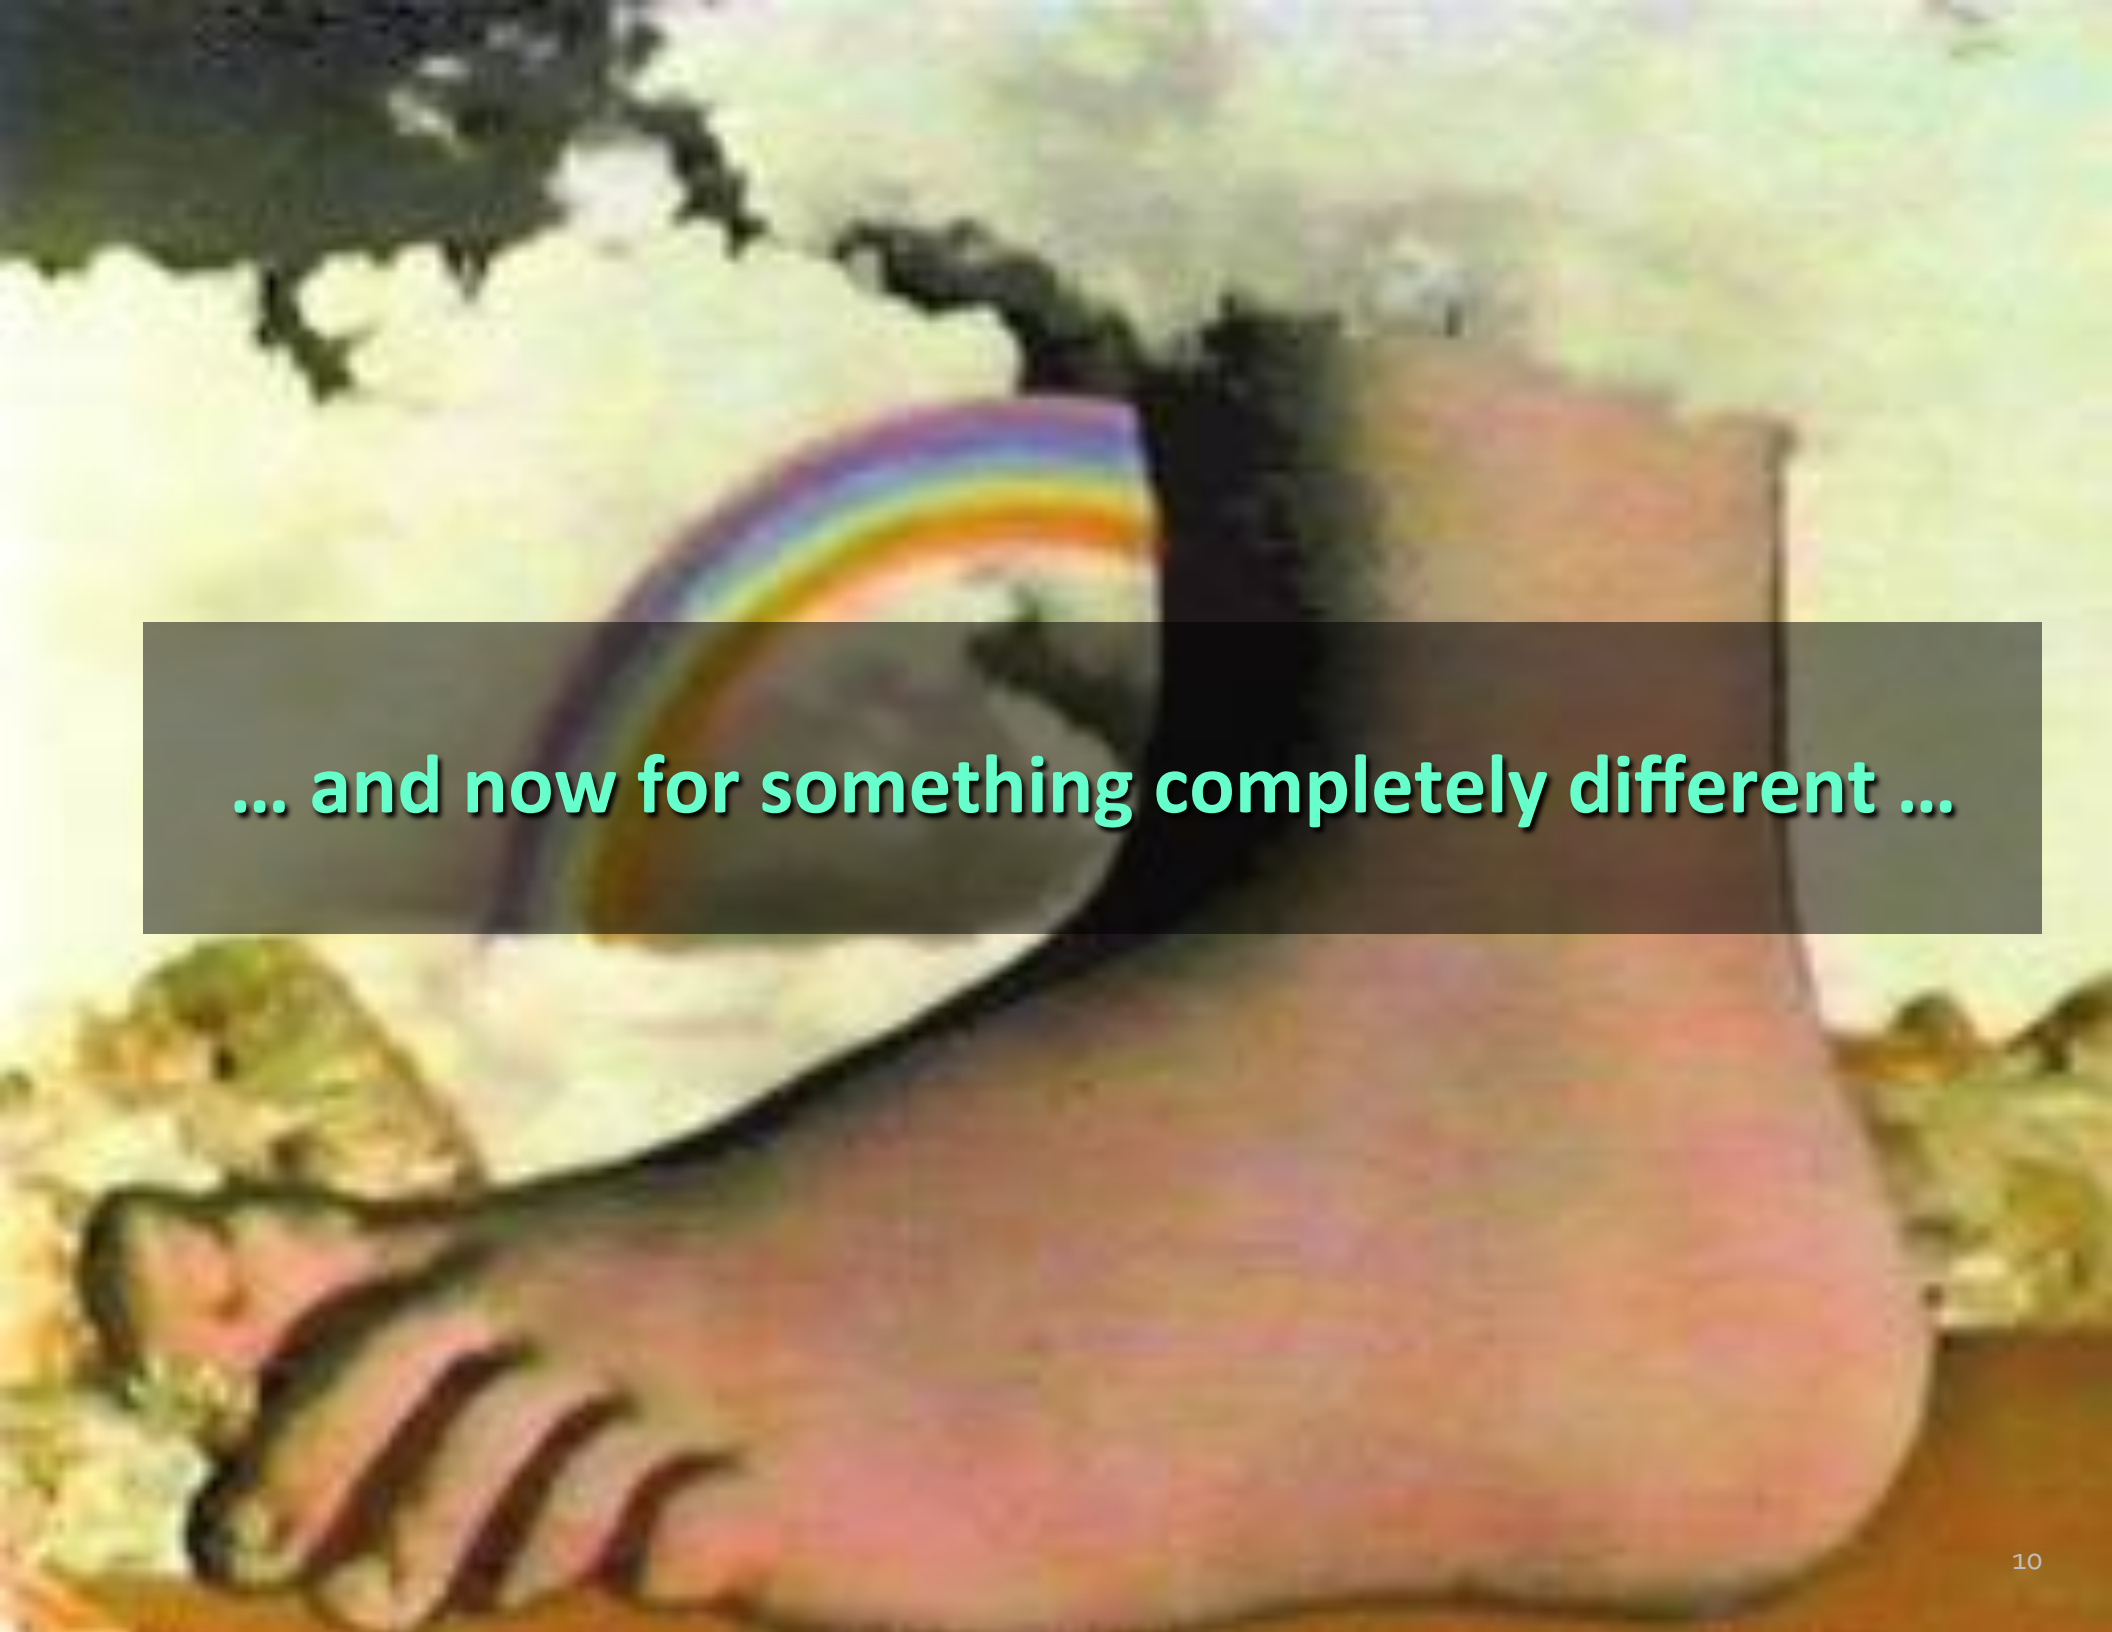
… and now for something completely different …

Elsewhere …
"big data" world is moving into RAM

- 2003: All Web indexes are served from RAM
- 2009: At Facebook 150 out of 200 TB cached
- New RAM data stores (***not*** caches)
  - Commercial: MemSQL, VoltDB, SQLFire,
  - Research: RAMCloud, H-Store, HyperDex

- General-purpose storage system
- All data always in DRAM (no cache misses)
- Durable and available
- Scale: 1000+ servers, 100+ TB
- Low latency: 5-10µs remote access

“Say ‘hello’ to my little friend”

- Dell PowerEdge R815
  - 48 cores @2.6GHz AMD
  - 256 GB RAM
  - 10Gb Ethernet
- Price?
  - 13-18 *iPads* !!

- 4 x R815 == neat little cluster:
  - 192 cores
  - **1TB** RAM

Fun things to do on 48 cores
(and 256GB of RAM)

- `pbzip2 -p48 target.dd` → 272MB/s
- `pbzip2 -d -p48 target.dd.bz2` → 677MB/s
- `pigz -p 48 target.dd` → 832MB/s

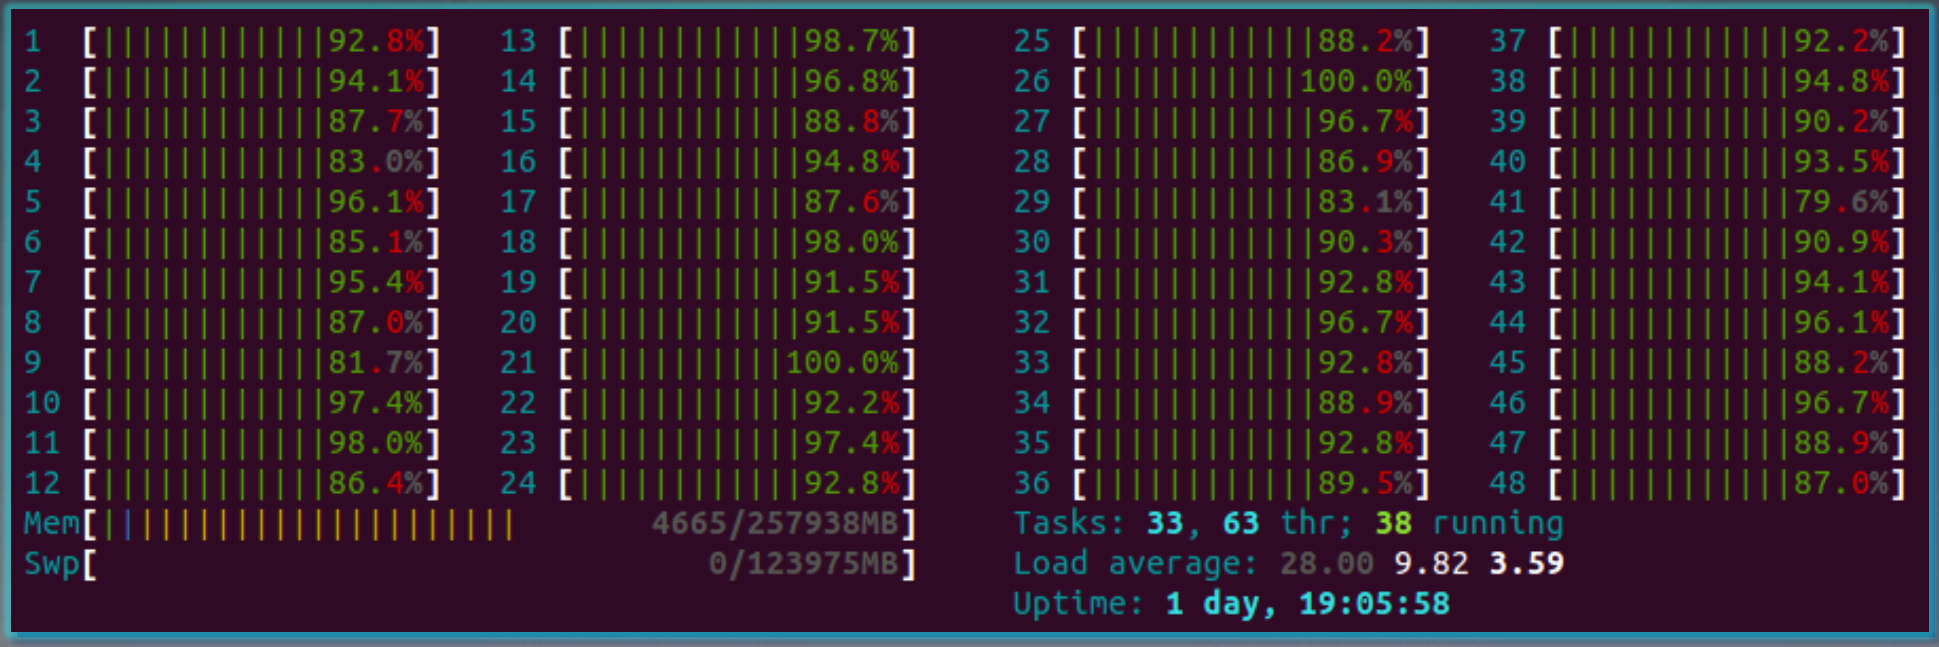

Unfun things to do on 48 cores
(and 256GB of RAM)

- `ewfacquire … target.dd` → 74MB/s
- `ewfexport … target.E01` → 147 MB/s

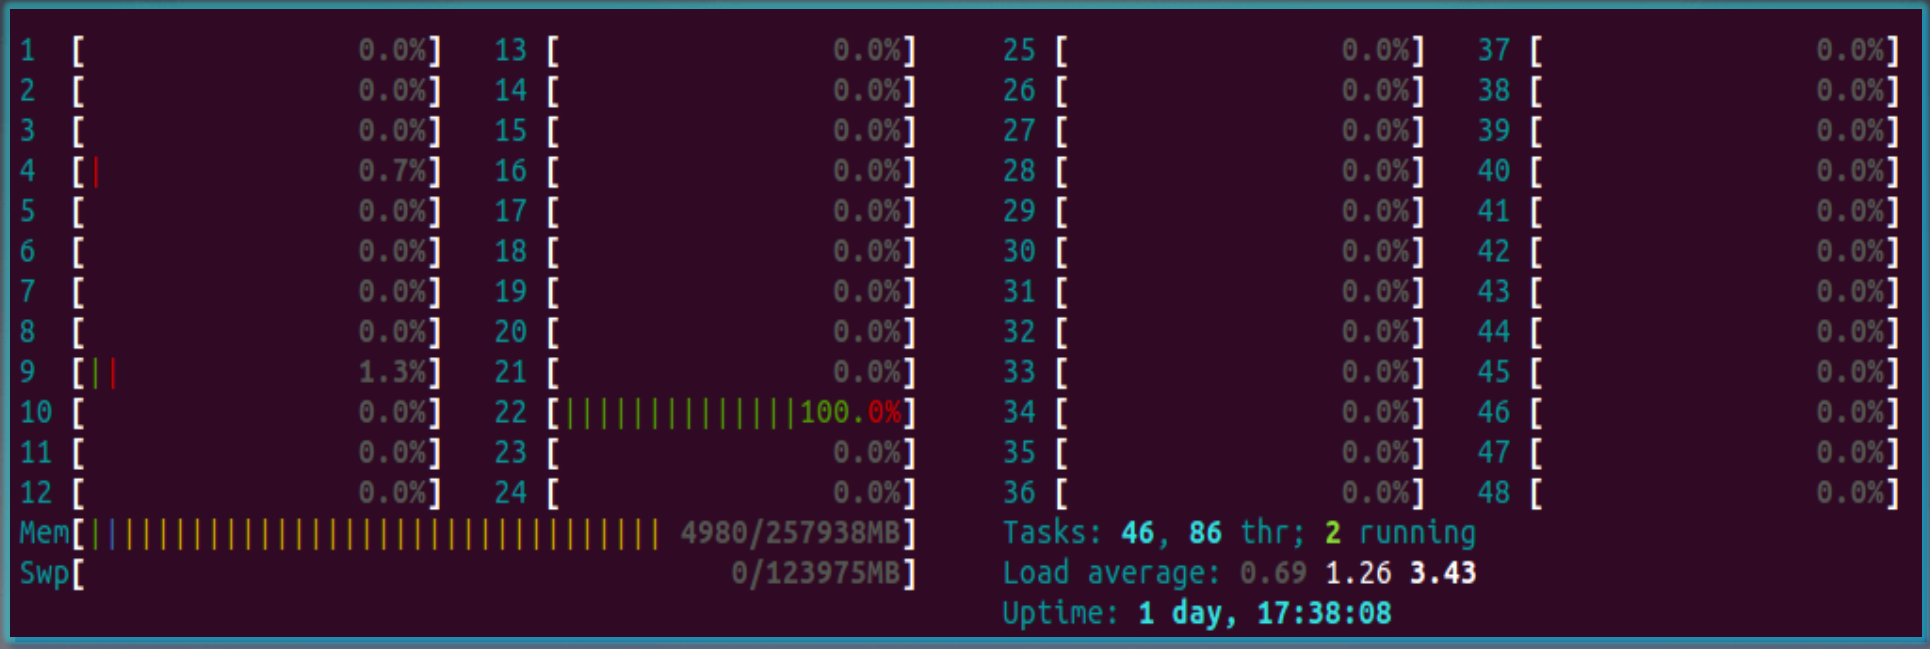

Useful things to do with 48 cores
(and 256GB of RAM)

- Screen content of a target *at line speed* with similarity digests (*sdhash 3.0alpha*):

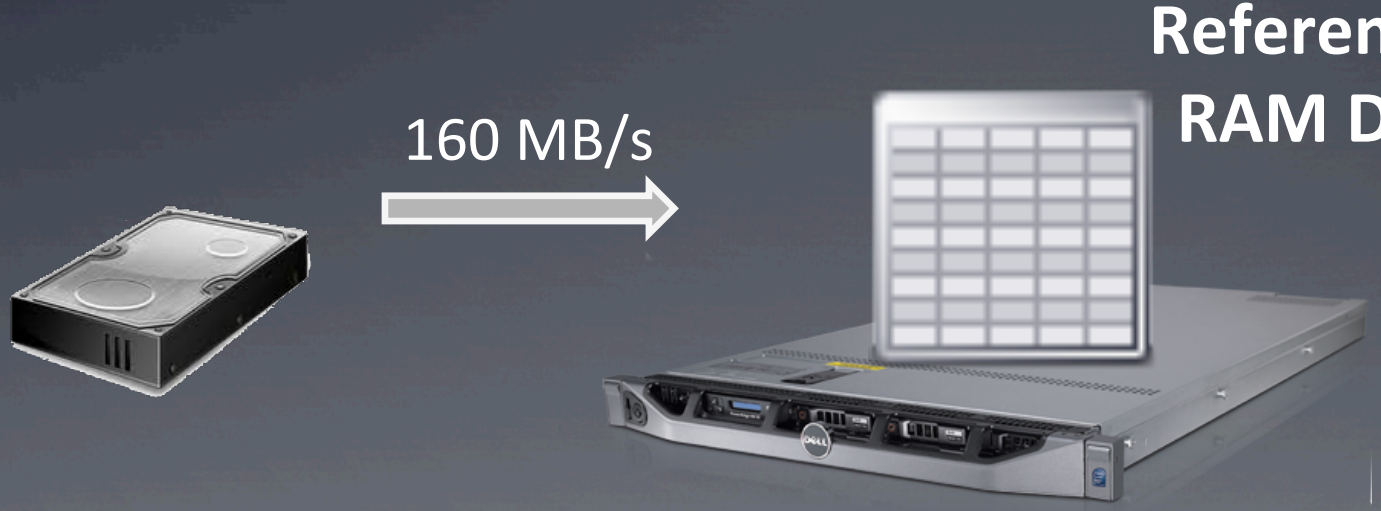

In V3.0 (Oct '12), we will cover up to 10GB of source data (RefDB: ~500MB)

In V3.1 (Dec'12), source data should be in the 1 to 10 TB range (RefDB: ~50-500GB)

Useful things to do with 48 cores
(and 256GB of RAM)

- Index (w/ Solr) 100GB in 40min
  - 43MB/s
  - Tested on GovDocs files: txt/html/doc/pdf/ppt/...
  - Zero disk I/O

- `time parallel exiftool -- /corpora/nps-gov/00?/*`

```
real    0m57.368s
user   37m41.317s
sys     2m52.163s
```

10,000 files, 5.5GB, cached

- `time parallel exiftool -- /corpora/nps-gov/01?/*`

```
real    1m37.142s
user   30m51.704s
sys     2m46.230s
```

10,000 files, 5.0GB, on disk

End goal: Real-time forensic processing

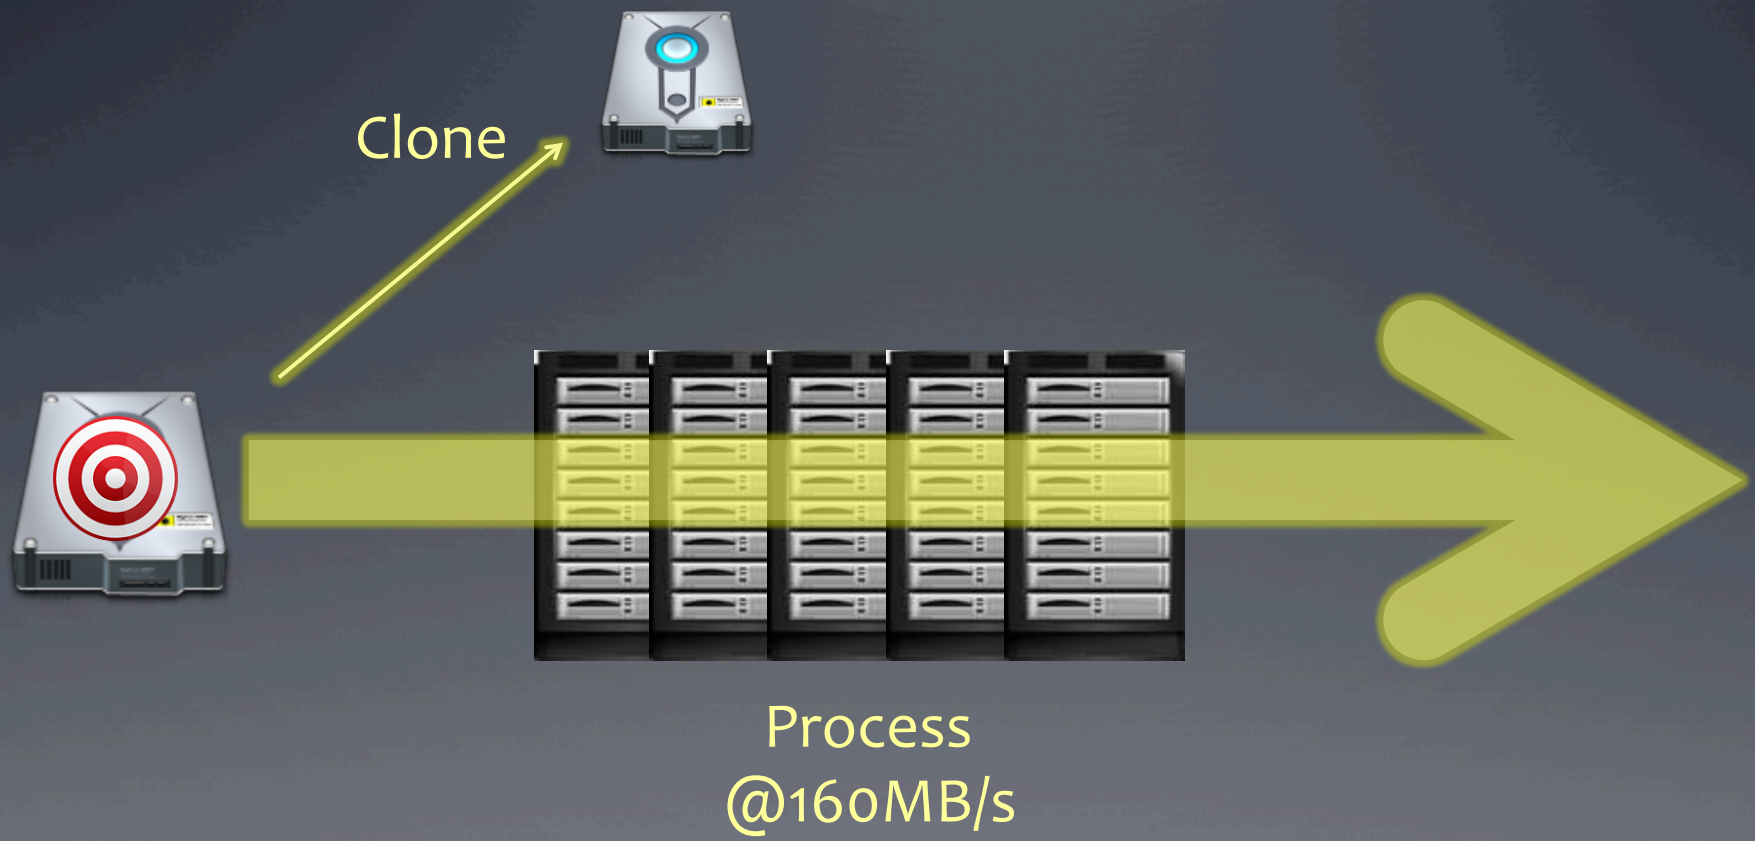

Objective: Finish cloning & processing at the same time.

Real-time forensics "showstoppers"

- **File**-based processing
  - hashing, metadata extraction, thumbnailing, …
  - generates non-sequential access & horrible I/O
- Indexing
  - search engines optimize query performance, not indexing
- Carving
  - can generate ***huge*** amounts of false positives (& potentially nasty I/O)

Latency-optimized target acquisition

- The problem
  - Most processing is file-centric
  - File-based access —> bad I/O on HDD
- Solution sketch
  - Before imaging, map blocks to files (45sec for 186,000 files)
  - During imaging, incrementally reconstruct files
    - Using multiple (potentially distributed) agents
  - Once file is complete, make it available via file system
  - File-based tool can pick it up and process it as usual
- End game
  - Given enough RAM/CPUs, time(cloning) == time(processing)

LOTA implementation

(by *Rob Martell*)

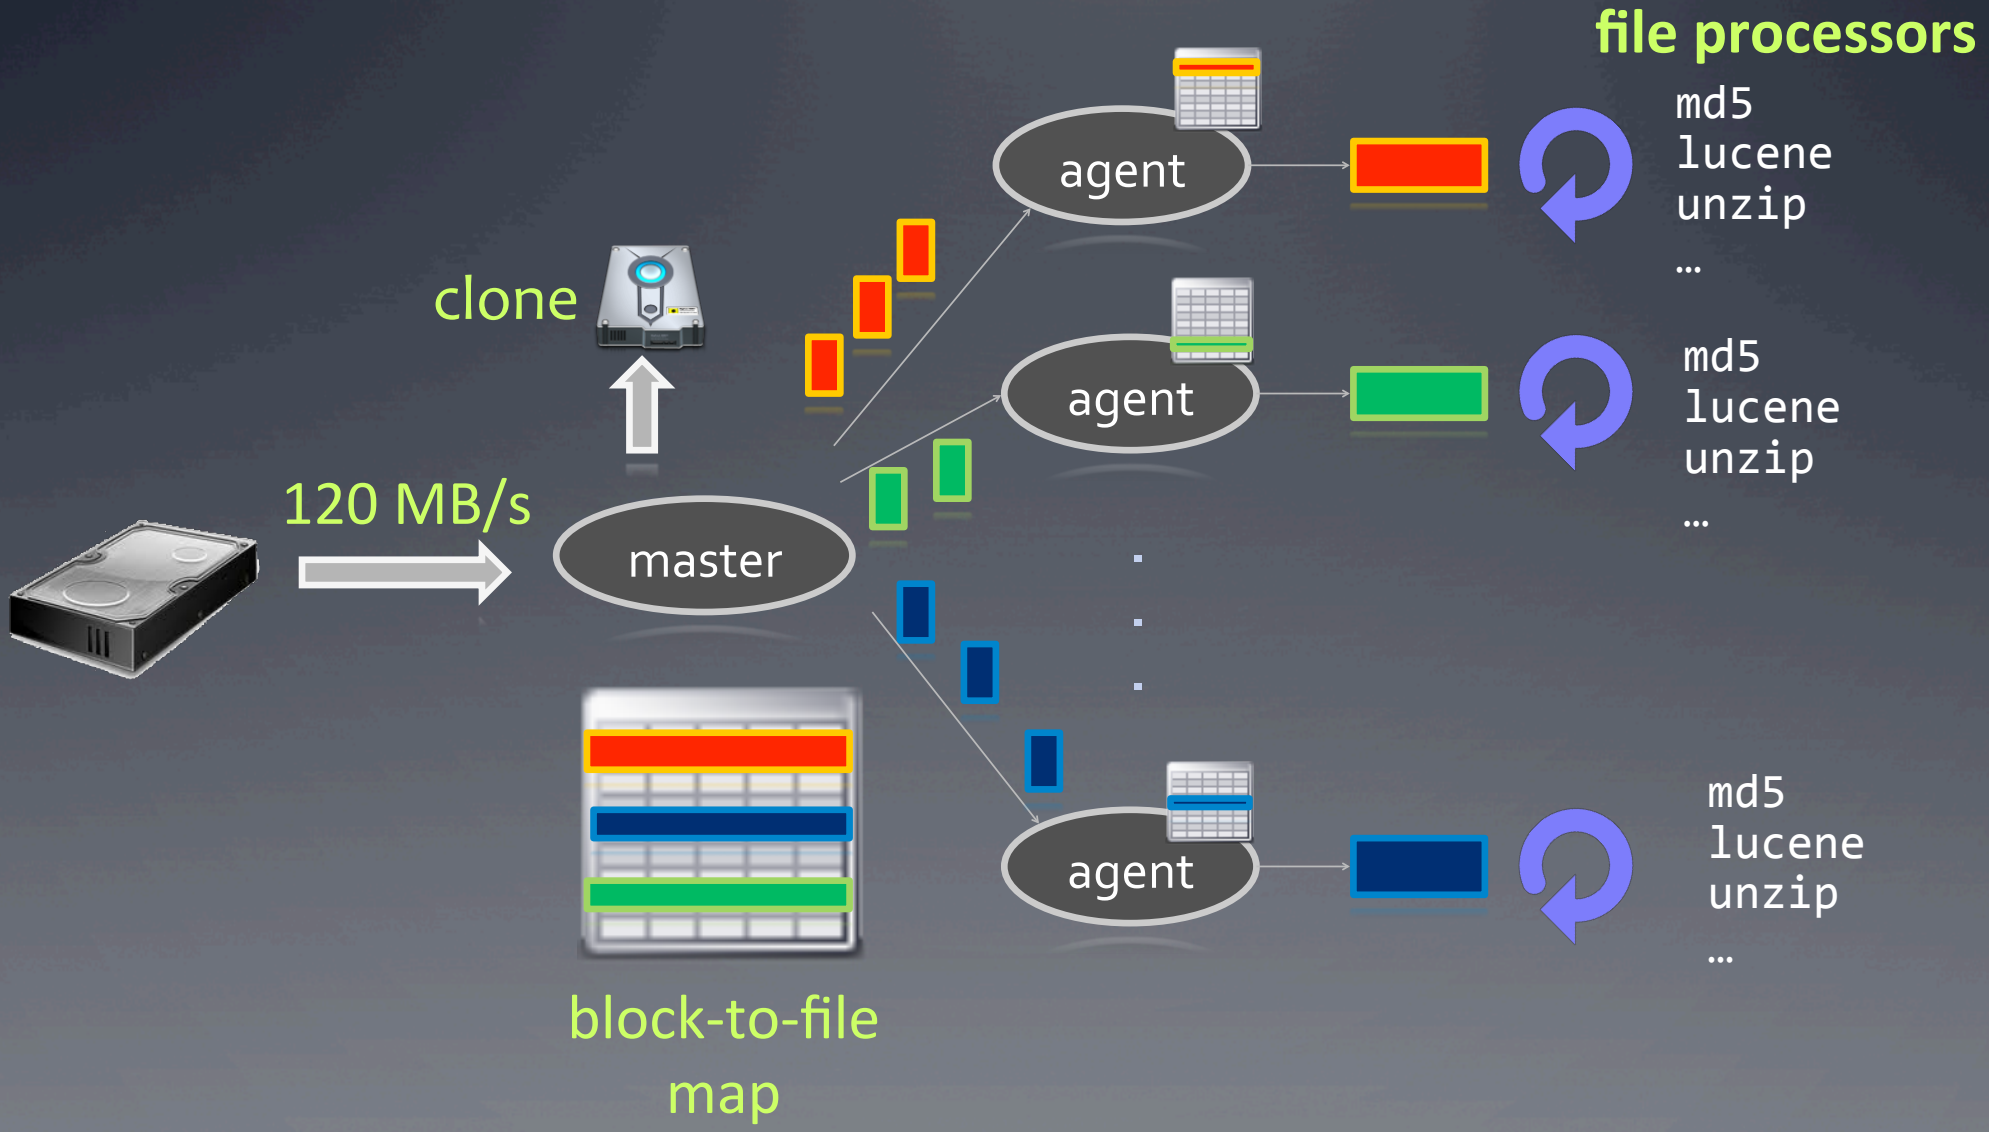

The takeaway (1)

- The primary performance concern of DF is ***latency***
  - Volume accumulation is a symptom of bad tools
  - Forensics needs to move to a real time model
    - ➔ **real-world processing has deadlines**
- Forensic analysis must start at conception
  - Move from 'clone-first' to 'latency-first' processing
- RAM will save the day (not *Hadoop*)
  - 10-1000x speedup for I/O-bound processing
  - Enables massive parallel processing

The takeaway (2)

- Current uses for (clusters of) high-RAM boxes
  - sdhash-based screening (NSRL scale)
  - latency optimized target acquisition (LOTA?)
  - indexing
  - metadata extraction, zip/unzip, thumbnailing
  - bulk_extractor
  - MySQL —> MemSQL
- We need a new platform to make all of this easy & extensible
  - We're working on it …

Thank You!

- Q & A
- Contact

  Vassil Roussev

  vassil@roussev.net
- sdhash 3.0 (exp. 10/15)

  sdhash.org