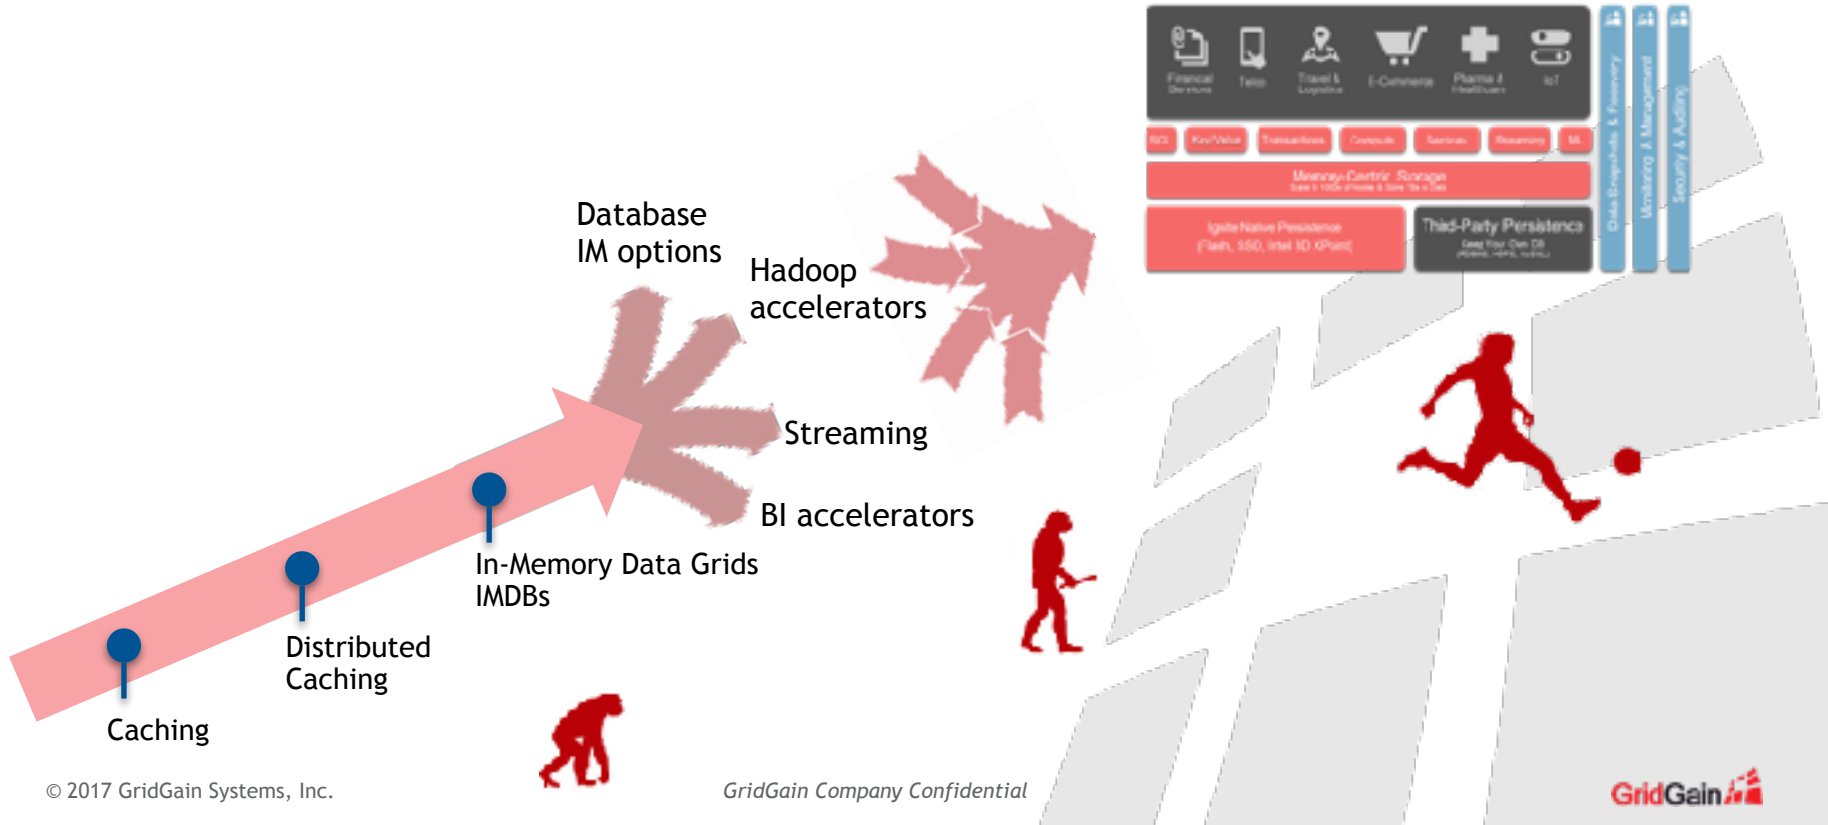




THE EVOLUTION OF IN-MEMORY COMPUTING PLATFORMS

Abe Kleinfeld
President & CEO

Evolution of In-Memory Computing



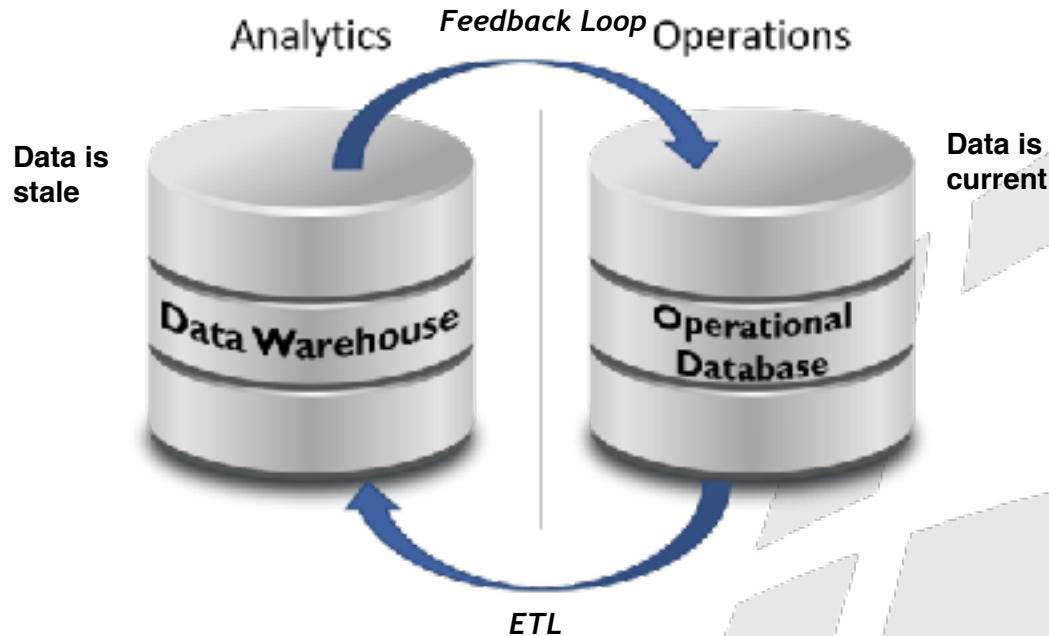
Disk-Based Computing Platforms

- Legacy IT infrastructure is based on a 50-year old model of data processing where data is stored and managed on disk
 - Optimized to read from / write to disk
 - Complex algorithms have emerged to improve performance for specific use cases, i.e. OLAP cubes, star and snowflake schemas
 - Caching and related techniques use memory to speed up specific functions, but **disk is the system of record**



Side-Effect of Disk-Based Computing

Decisions are made on yesterday's information



In-Memory Computing Emerges

- In-Memory Computing has emerged due to the collision of massive **data growth** with **internet-scale**, resulting in an insatiable need for speed and scale



In-Memory Computing Platform


- The new IT infrastructure where data is stored in memory
 - Optimized for to read from / write to memory
 - Entire data set in stored in memory to maximize speed
 - Employs distributed architecture for horizontal scalability
 - Disk is used for backup/persistence, but **memory is the system of record**



In-Memory Enables the Real-Time Enterprise (Hybrid Transaction/Analytical Processing)

“TMC-enabled HTAP can have a transformational impact on the business.” — Gartner 2/17

Disk-Based




In-Memory HTAP Architecture



HTAP enables real-time analytics and situation awareness on live transaction data as opposed to after-the-fact analysis on stale data (in traditional architectures).

But the truth is...

- IMC must evolve to embrace data growth cost effectively
 - Through the intelligent use of new memory technologies
 - And yes, the use of SSD/disk



The Next Step in IMC Evolution

- Process massive data sets
- Deliver IMC performance
- Meet scale expectations
- Offer more granularity of performance and cost tradeoffs
- Embrace new storage-class memory technologies



Memory-Centric Computing

The Next Logical Step for In-Memory Computing

- Memory is still the system of record
- Datasets can exceed available RAM by seamlessly extending to non-volatile, storage-class memory, including disk

Disk-based systems use **memory to cache** frequently used data

Memory-centric systems use **disk to stash** less used data

Memory-Centric

Data in RAM

Speed



Cost



Memory-Centric

Data in RAM

Speed



Data in NV Memory

Cost



Memory-Centric

Data in RAM

Speed

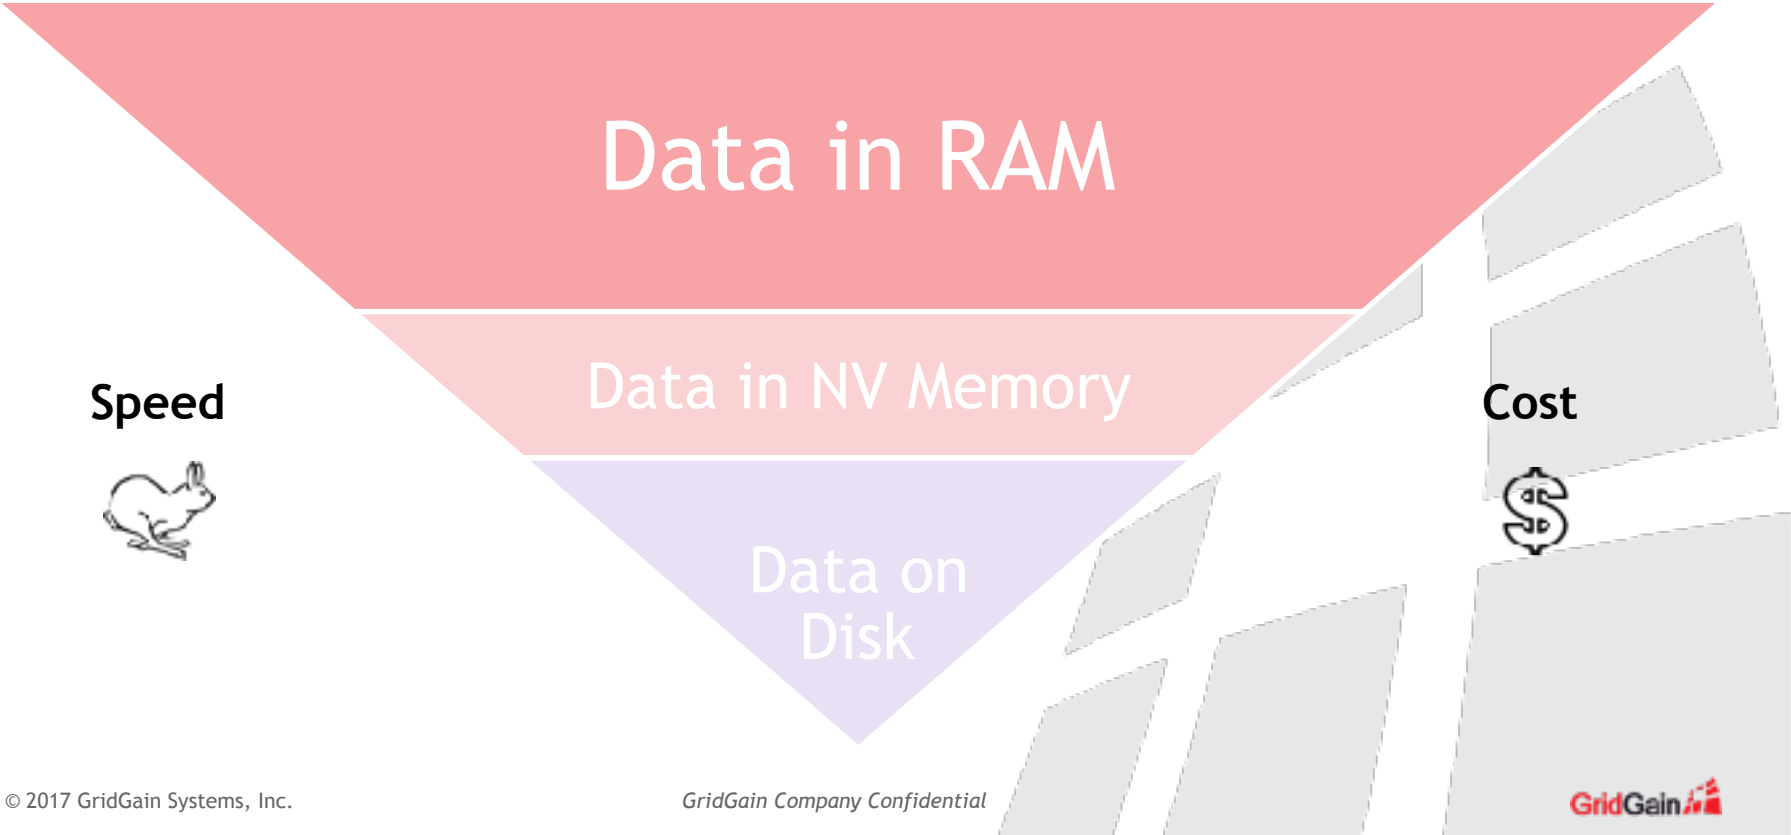


NVDIMM ReRAM
PCM NRAM
3D XPoint


Cost



Memory-Centric




Memory-Centric




Why Memory Centric Computing?

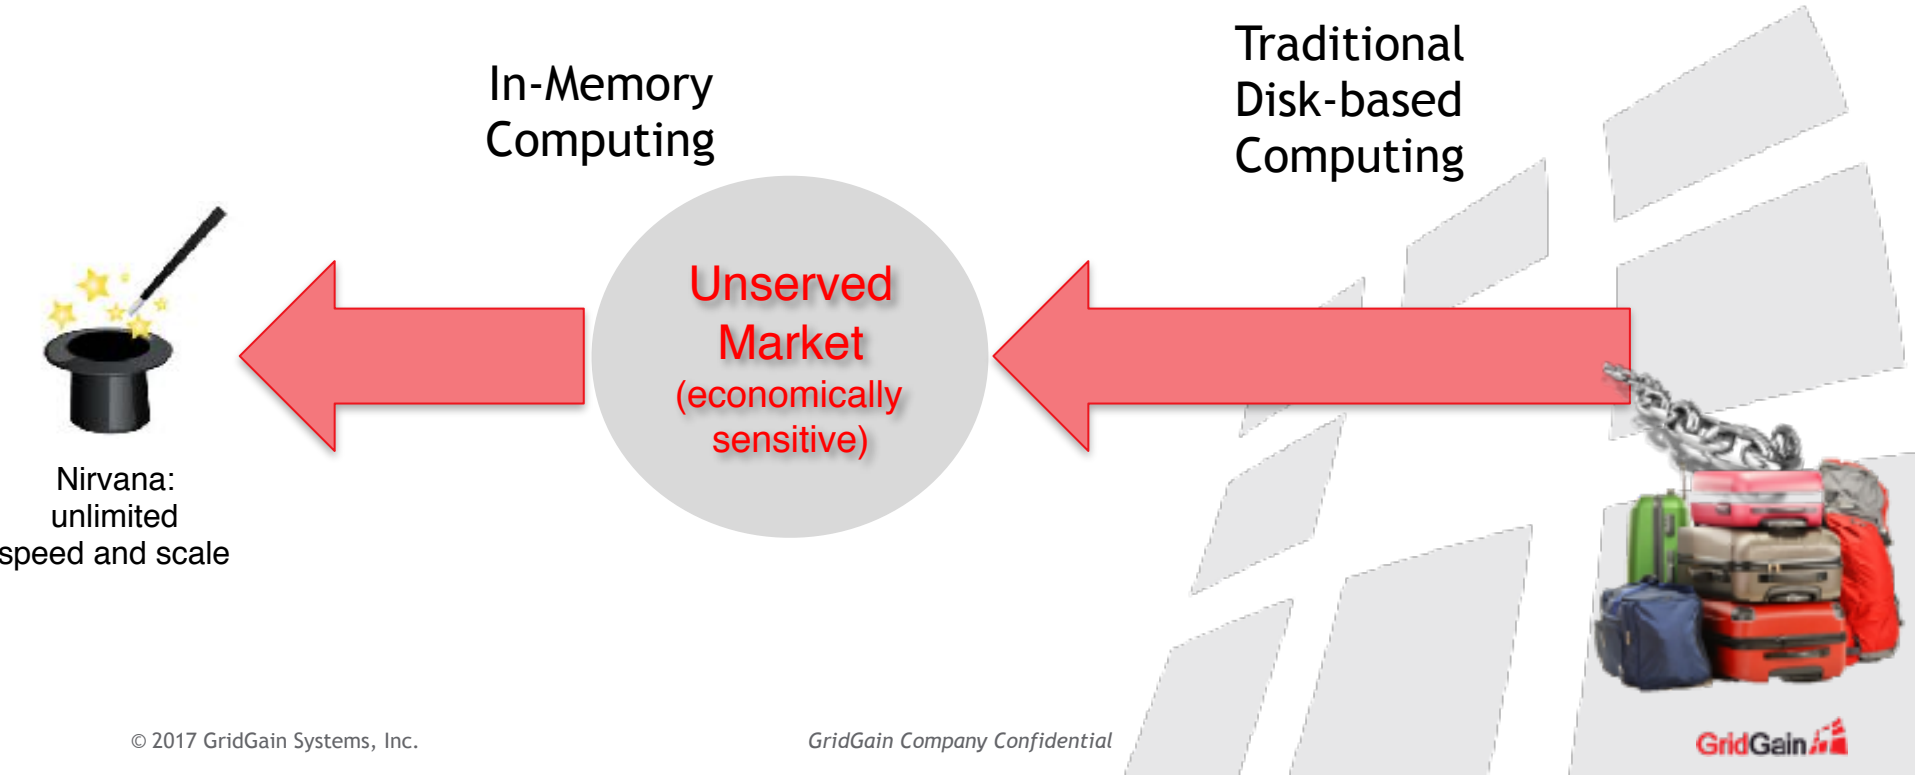
- Benefit of IMC speed with ability to balance performance and economics
- Memory centric architectures are built with a memory-first approach, without the legacy baggage of disk-first systems
- Modern memory centric architectures are naturally distributed to easily scale out (and up)
- No memory warm-up required on restarts in memory-centric computing



Convergence of in-memory and disk computing?



Race to Nirvana




Race to Nirvana

Memory-Centric
Computing

Traditional
Disk-based
Computing



Nirvana:
unlimited
speed and scale



Race to Nirvana

Memory-Centric
Computing

Traditional
Disk-based
Computing



Nirvana:
unlimited
speed and scale



Summary

- IMC platforms are evolving to become
- No reason to build greenfield applications on IMC platforms anymore
- MC can deliver everything IMC delivered at better economics
- Data set size is no longer a limitation
- For legacy application modernization, migrate to cache disk-based platforms, and eventually to a cloud-based disk-based platform

