SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Optimizing Complex Memory and Storage Systems using Simulations

Presented by Andy Banta, Magnition IO

Andy Banta

Magnition.io (Consultant)
SolidFire (VMware development, acq. by NetApp)
DataGravity (Container exploitation lead)
VMware (iSCSI Tech Lead, IPO)
Sun Microsystems (Initial Fibre Channel development)
Patent, early distributed network projects, data acquisition
@andybanta

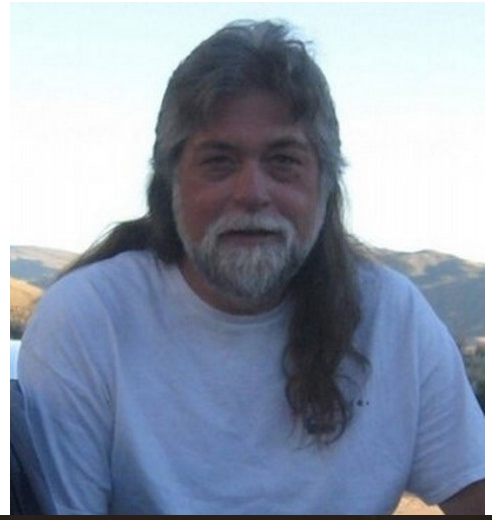

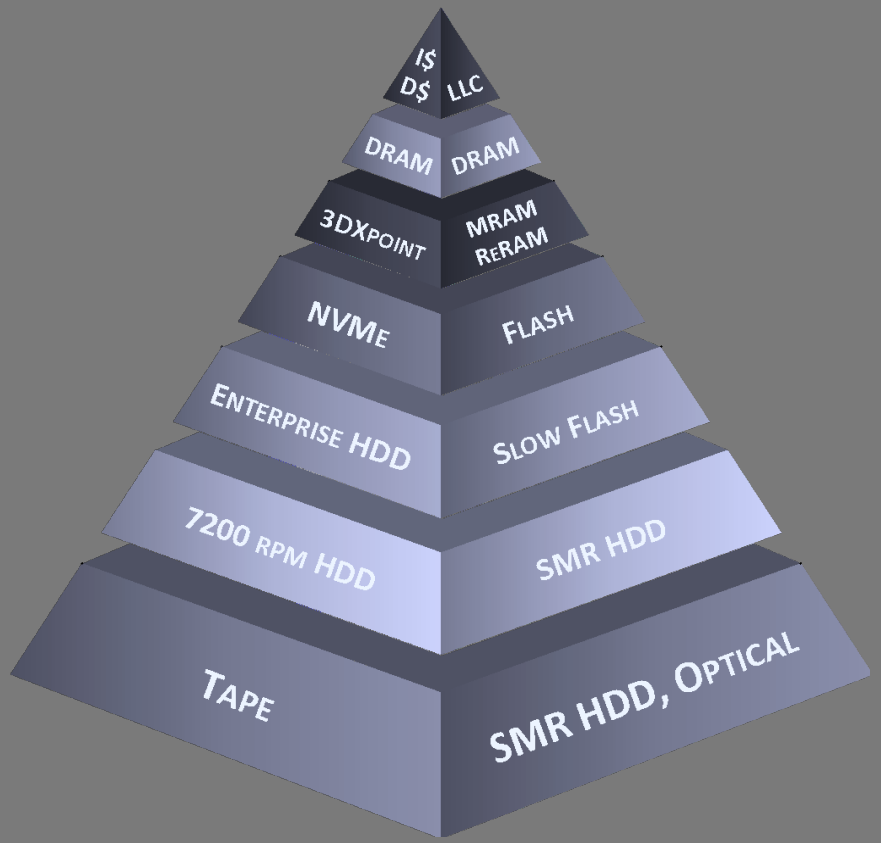

The Challenge

Modern compute and storage system use multiple layers interacting in multiple ways

HOW CAN CURRENT TECHNOLOGY ACHIEVE…

- Latency control
- Multi-tenant thrash remediation
- Correct tier sizing
- Workload-awareness
- Hot working set management
- Latency and throughput SLAs
- Memory capacity planning

AS MORE HARDWARE LAYERS ADD COMPLEXITY?

ABOUT MAGNITION

STORAGE PERFORMANCE, REINVENTED

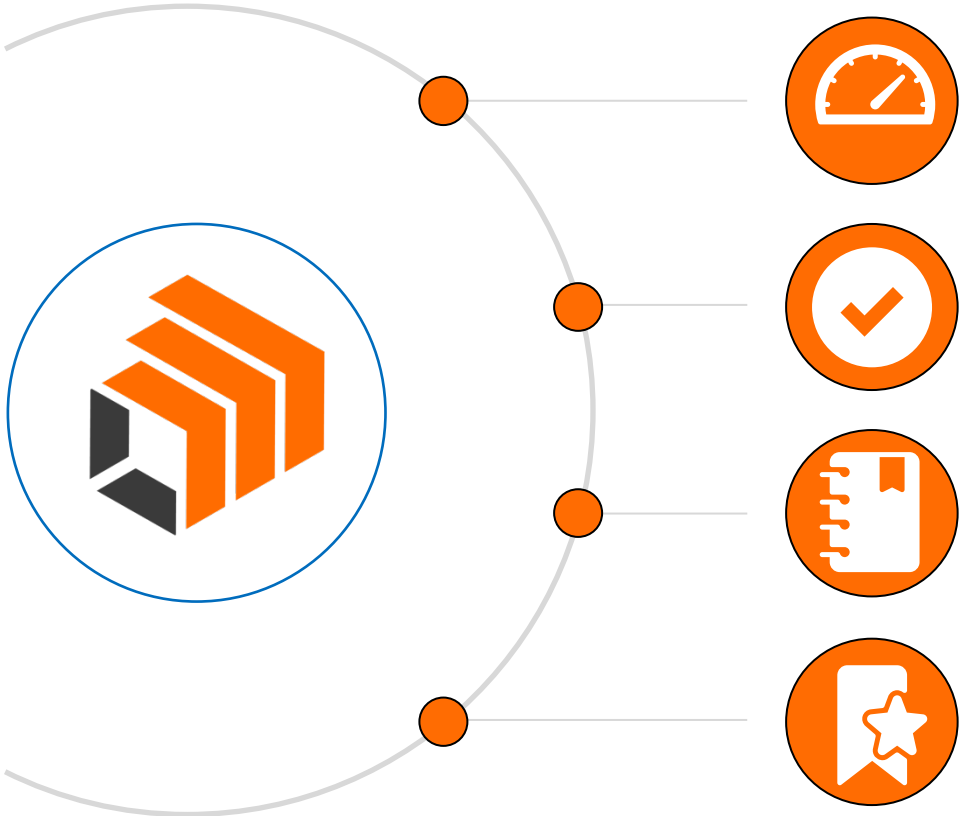

World's First Real-Time Data Placement Optimization

Patented technology is a first for the industry.

Proven At-Scale, with Production Workloads

Use customer traces to fully test diverse workloads in real-time.

Peer-Reviewed and Published in Leading Journals

Multiple industry articles published and reviewed.

Award-Winning, Patented Technology

3-time award winner for innovative technology.

A different approach to optimization

- Compose simulations of complex memory and storage
- Break the simulation into components
- Allows the components to be assembled like building blocks
- Provide reasonable but constrained set of variables
- Run simulations with synthetic data or actual IO traces

Value of simulations

- Faster and easier to prototype
- Minimal up-front hardware spend
- Great opportunities for optimizations
- Loads of simulations are done at ASIC level
  - The same practices should apply to component and software levels
- Choose three
  1. Lower cost
  2. Higher speed
  3. More flexibility

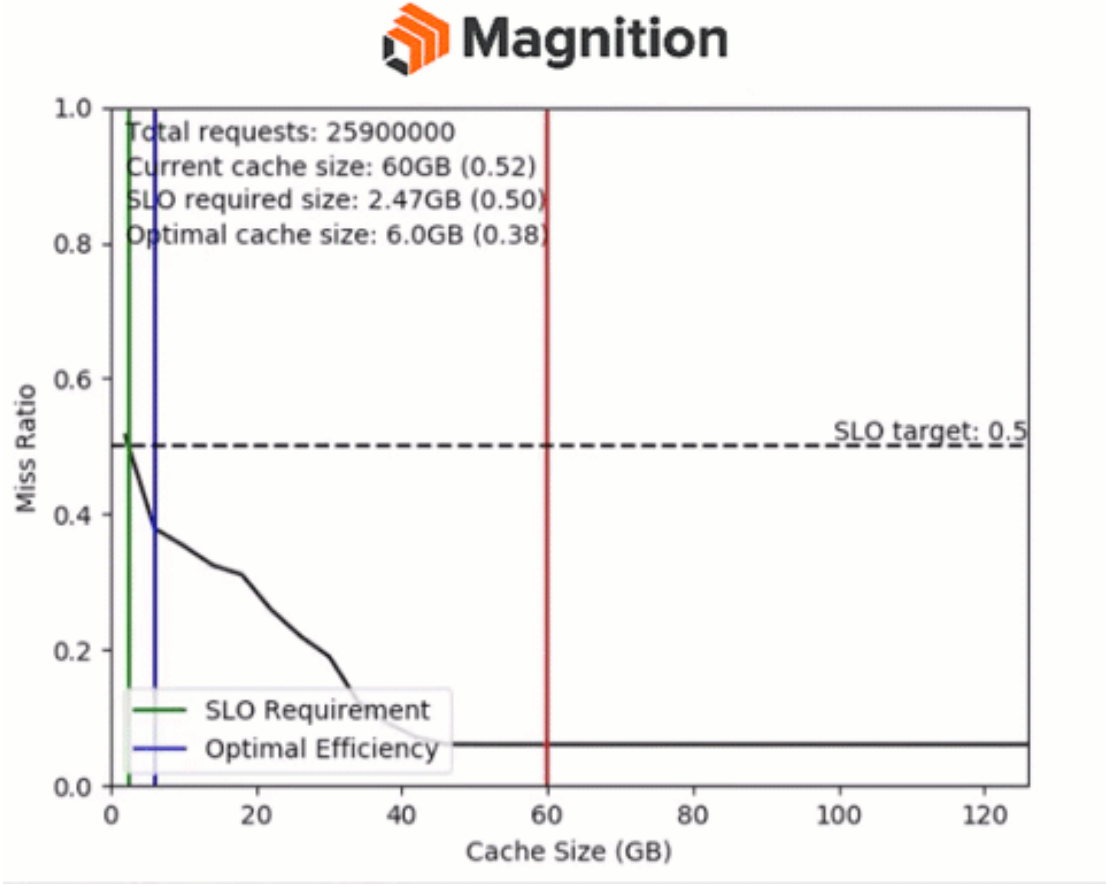

Composable components

- Provide a framework to connect components
  - Lingua Franca provides this
  - Reactors represent system pieces
- Library of ready components for use
- Allows clients to build their own modules
- Basic set of building blocks
  - Cache
  - Media
  - Wire

Composable components

- Provide a framework to connect components
  - Lingua Franca provides this
  - Reactors represent system pieces
- Library of ready components for use
- Allows clients to build their own modules
- Basic set of building blocks
  - Cache
  - Media
  - Wire

Composable components

- Provide a framework to connect components
  - Lingua Franca provides this
  - Reactors represent system pieces
- Library of ready components for use
- Allows clients to build their own modules
- Basic set of building blocks
  - Cache
  - Media
  - Wire

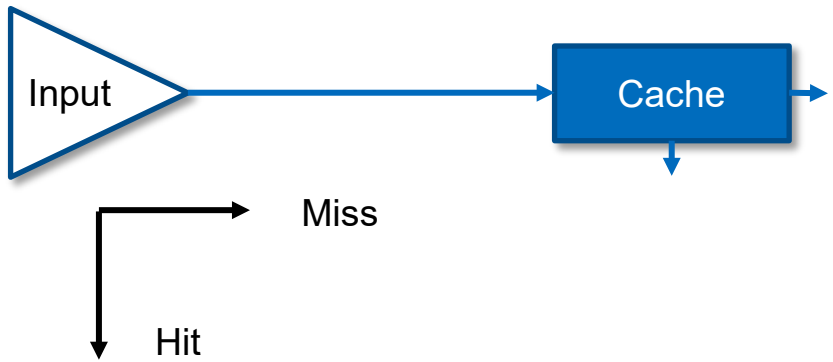

Composable components

- Provide a framework to connect components
  - Lingua Franca provides this
  - Reactors represent system pieces
- Library of ready components for use
- Allows clients to build their own modules
- Basic set of building blocks
  - Cache
  - Media
  - Wire

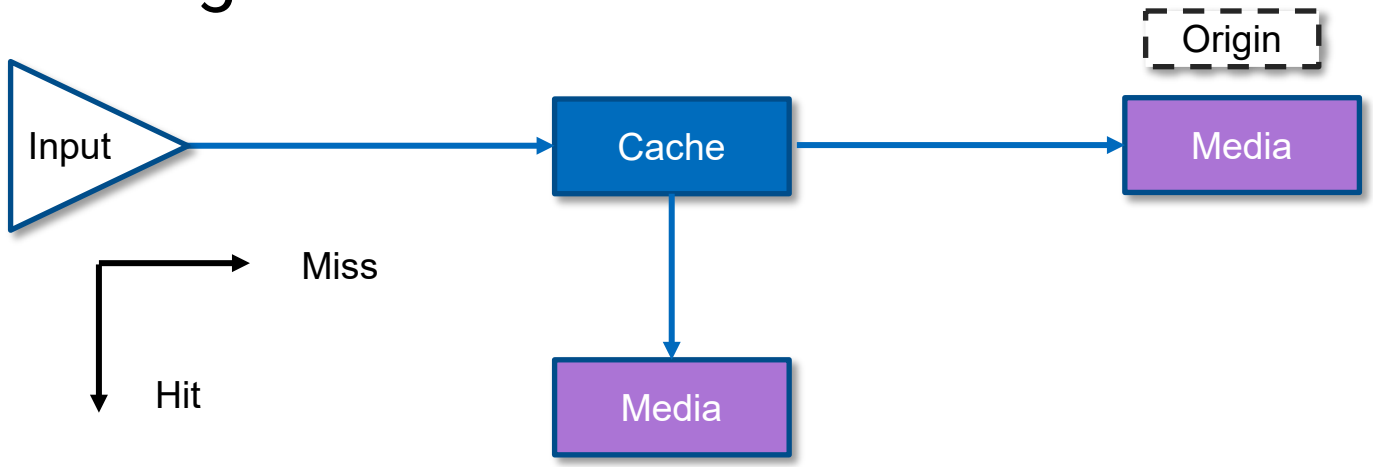

Composable components

- Provide a framework to connect components
  - Lingua Franca provides this
  - Reactors represent system pieces
- Library of ready components for use
- Allows clients to build their own modules
- Basic set of building blocks
  - Cache
  - Media
  - Wire

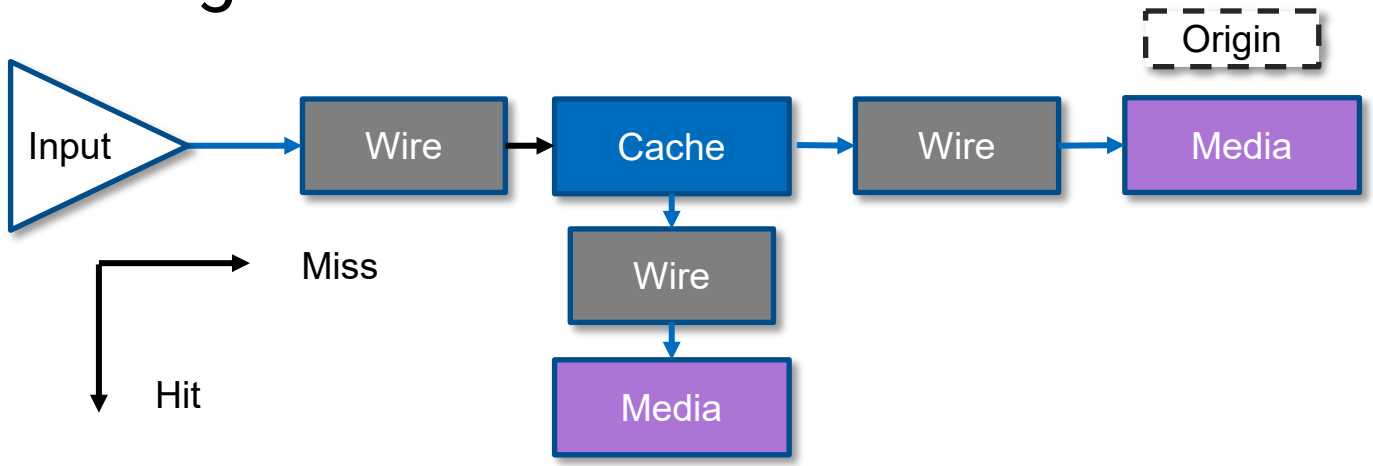

Composable components

- Provide a framework to connect components
  - Lingua Franca provides this
  - Reactors represent system pieces
- Library of ready components for use
- Allows clients to build their own modules
- Basic set of building blocks
  - Cache
  - Media
  - Wire

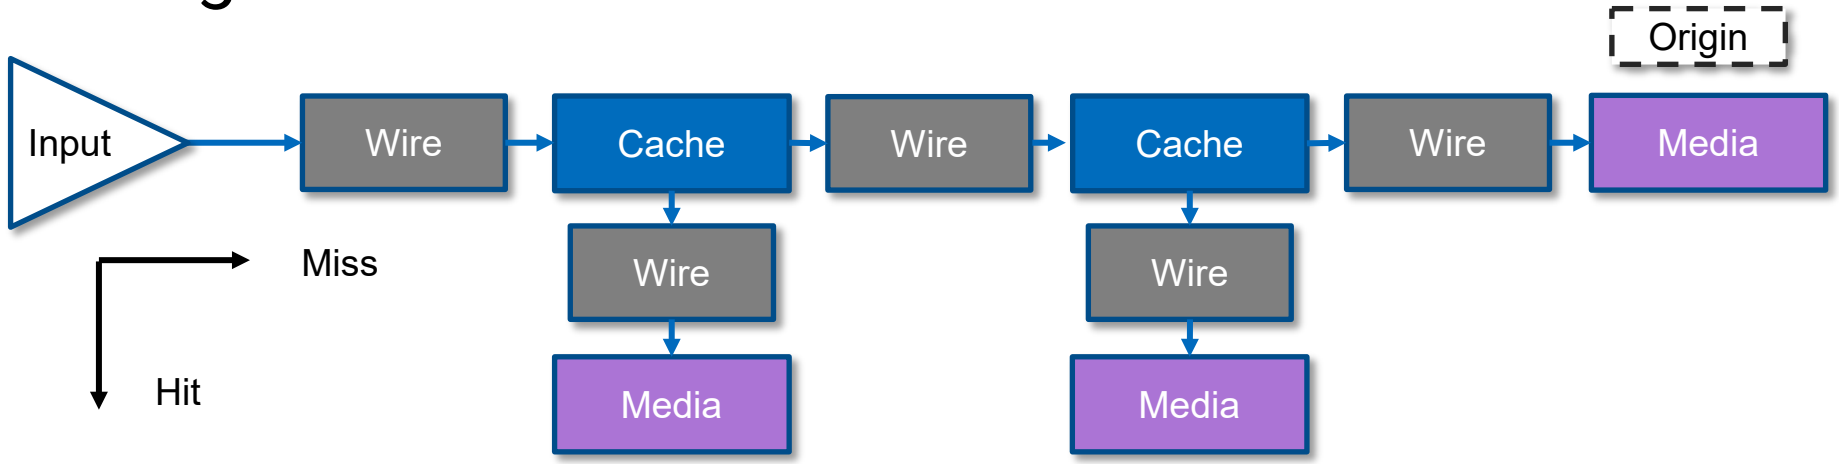

Cache component

- Easily build basics like lookups, allocation, and eviction
- One (or more) hit path
- One (or more) miss path
- Many choices for variability

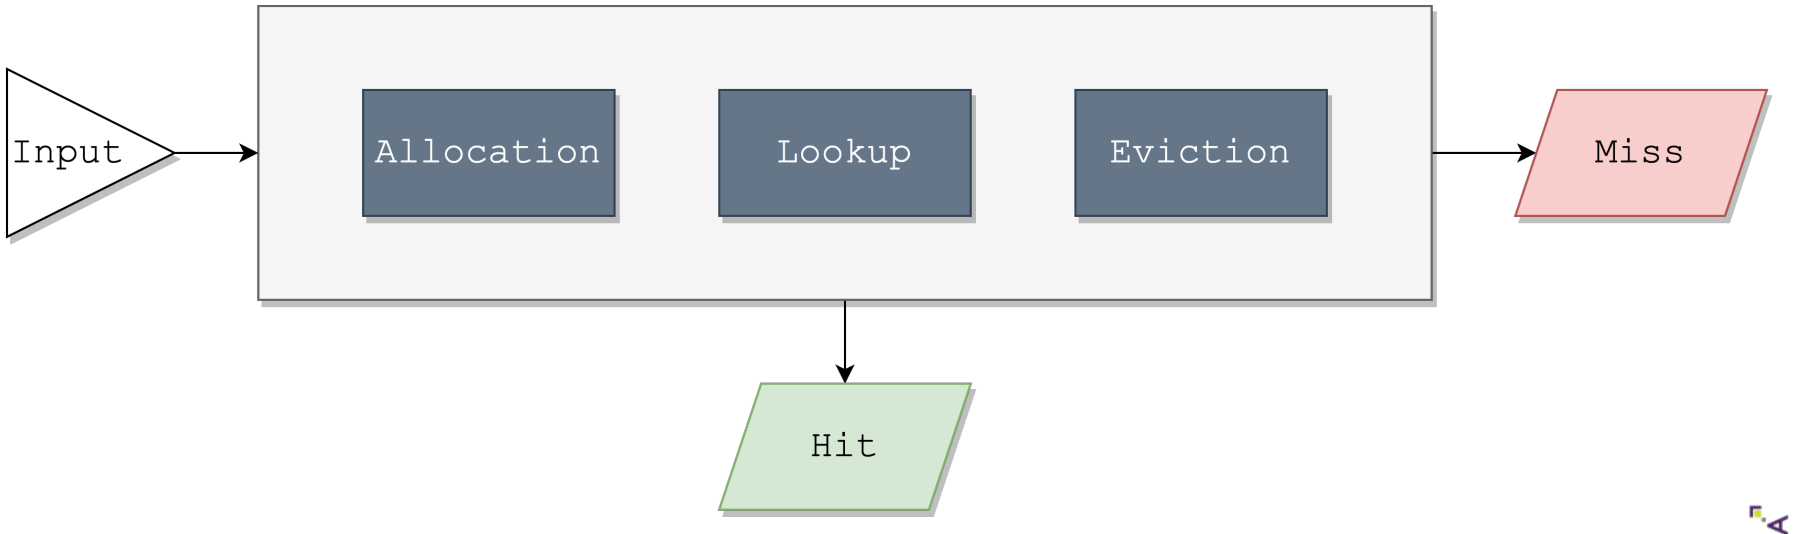

Media component

- Memory, disk, cloud storage
- Introduce distinct delays
  - MQSim
- Parallel access
  - Contention delays
  - Queueing
- Only need to simulate delay
  - Not actual data delivery

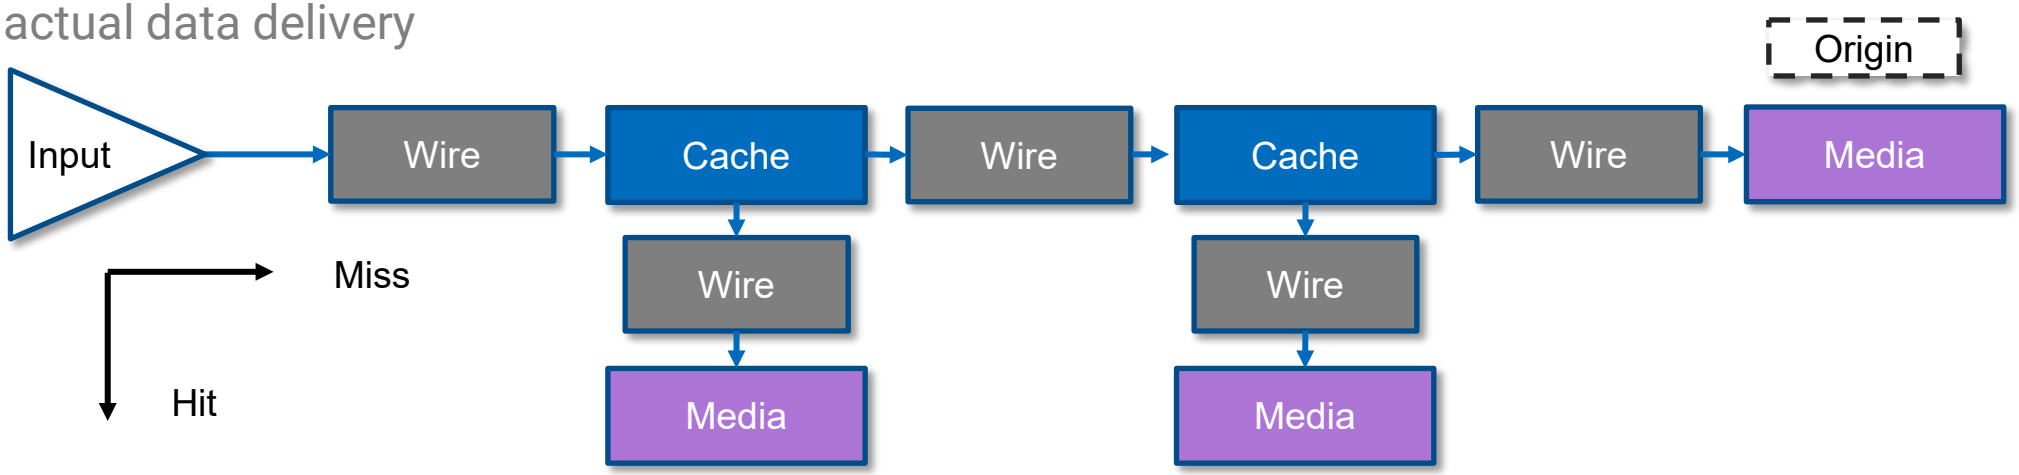

Media component

- Memory, disk, cloud storage
- Introduce distinct delays
  - MQSim
- Parallel access
  - Contention delays
  - Queueing
- Only need to simulate delay
  - Not actual data delivery

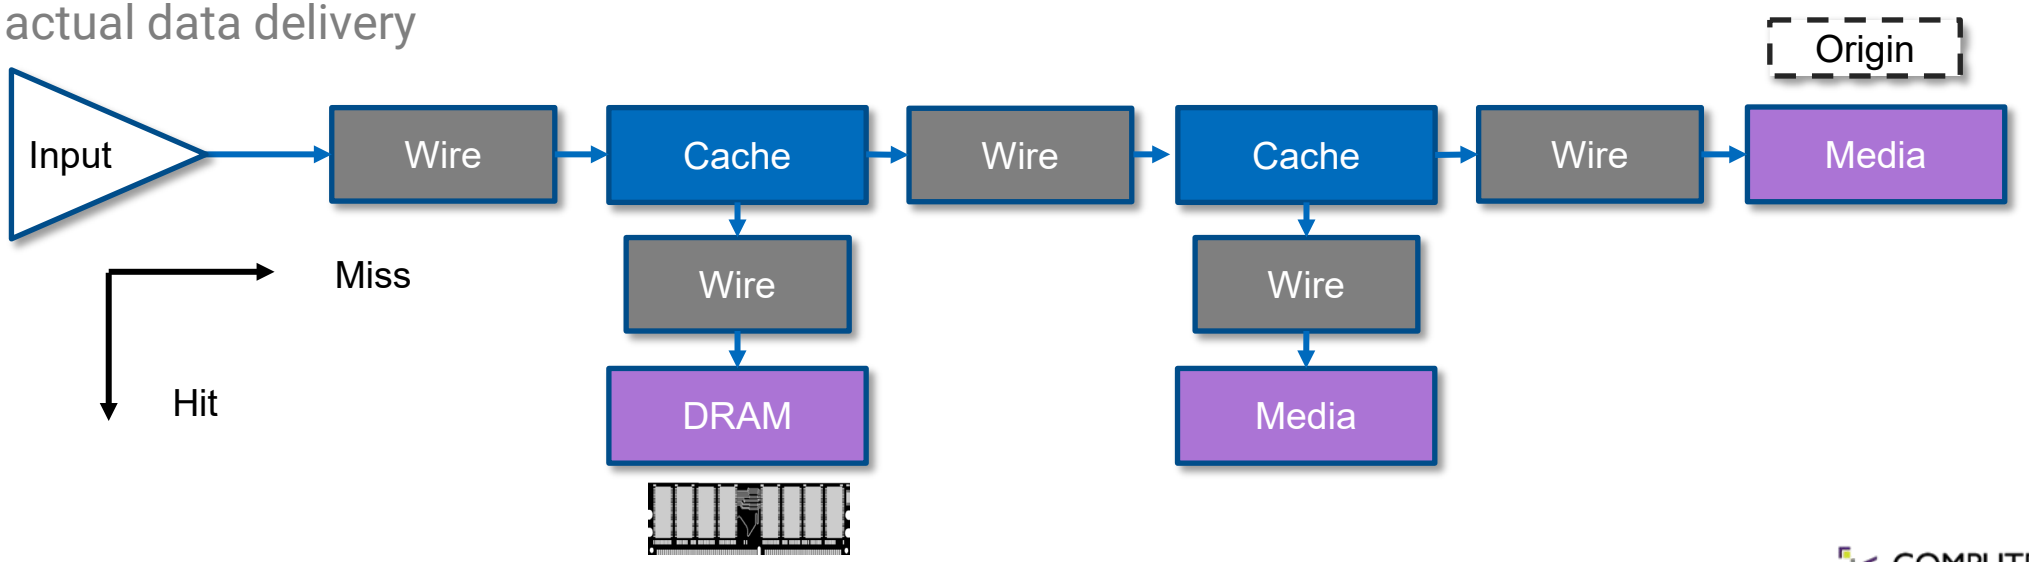

Media component

- Memory, disk, cloud storage
- Introduce distinct delays
  - MQSim
- Parallel access
  - Contention delays
  - Queueing
- Only need to simulate delay
  - Not actual data delivery

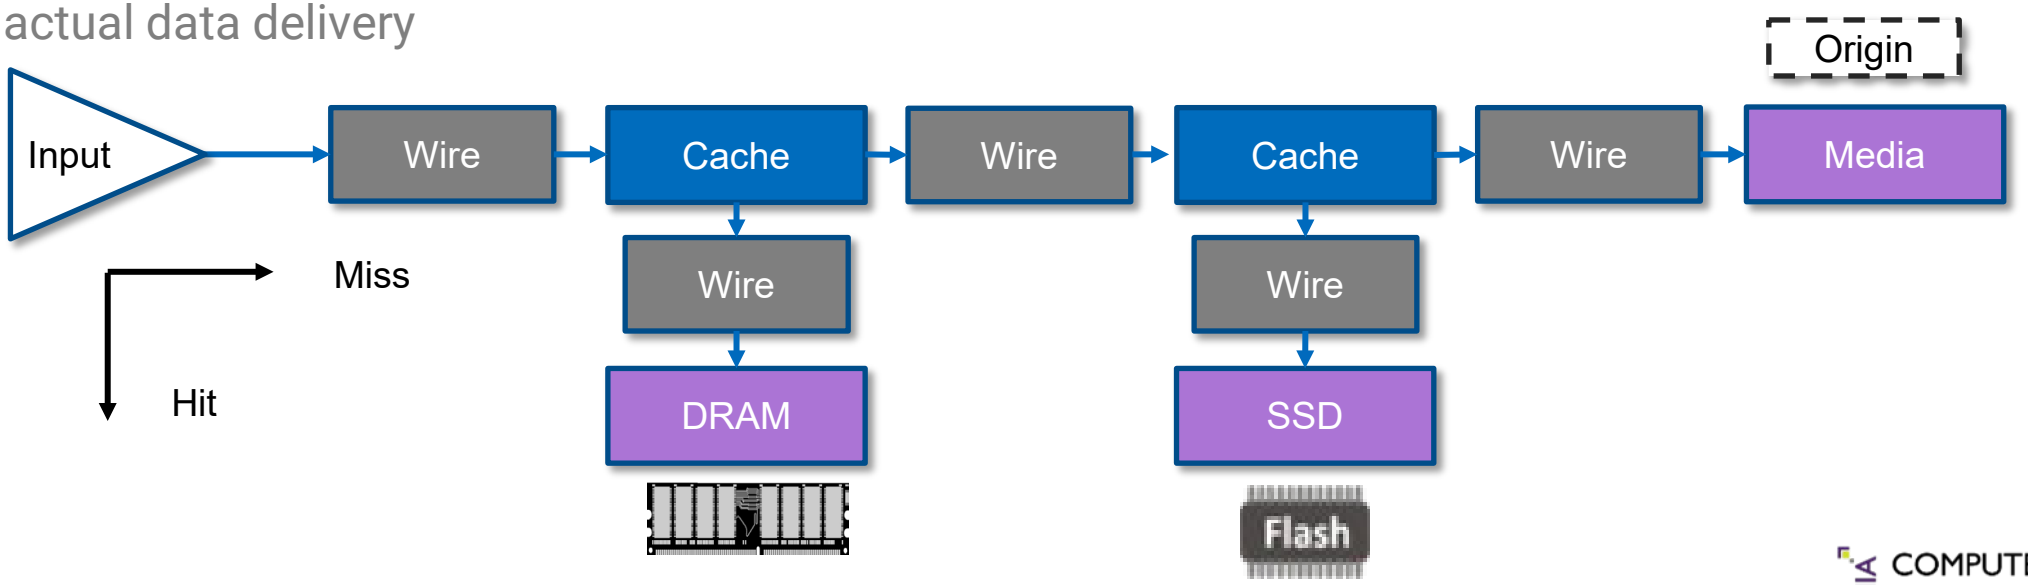

Media component

- Memory, disk, cloud storage
- Introduce distinct delays
  - MQSim
- Parallel access
  - Contention delays
  - Queueing
- Only need to simulate delay
  - Not actual data delivery

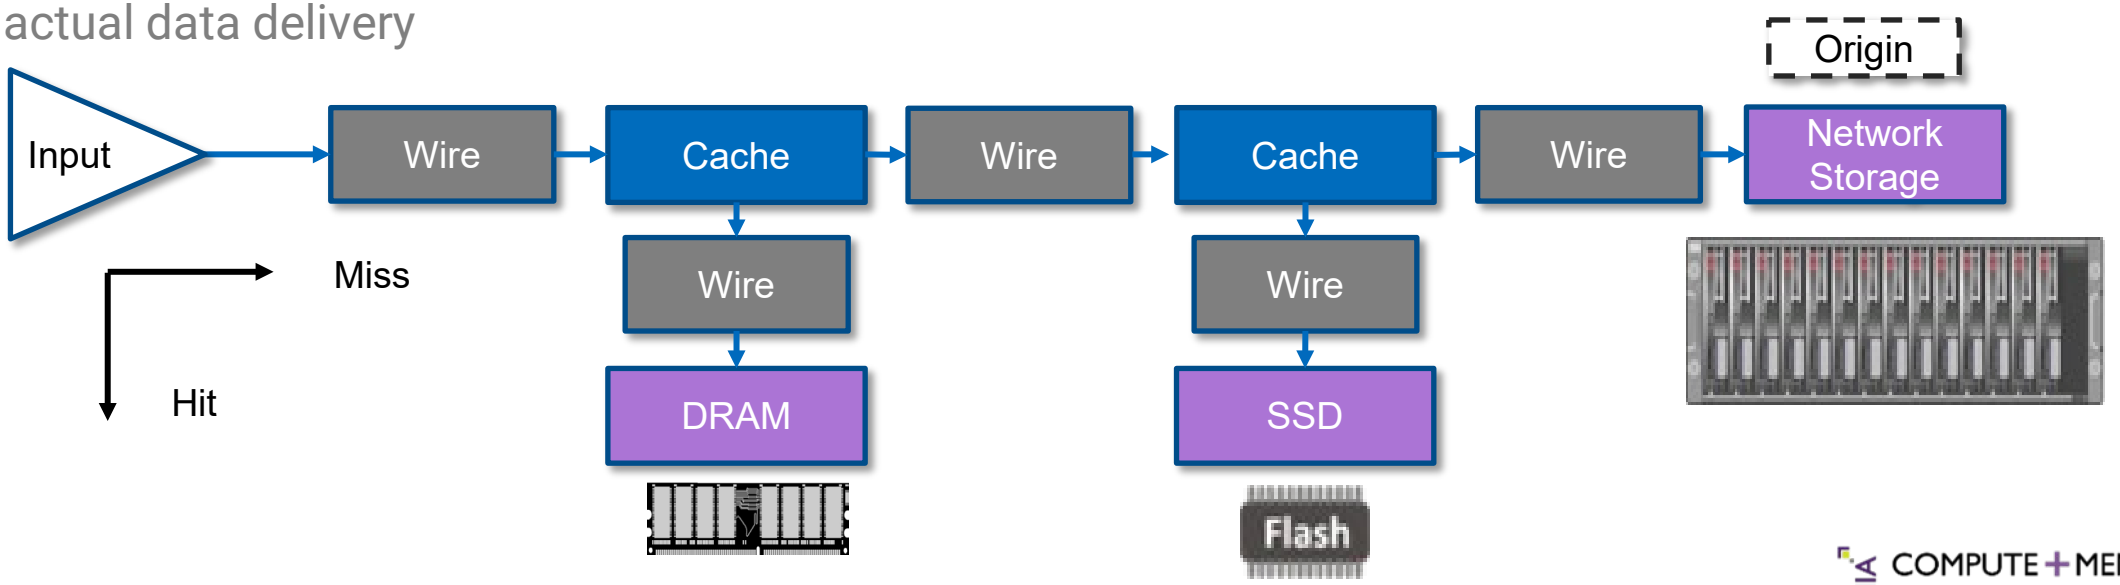

Wire component

- Memory bus, disk controller, network
- Can multiplex and change form of IO request
- Even type of wire can be variable
  - Type of memory bus
  - Hops in network topology
- Delays introduced by wire, contention, queueing

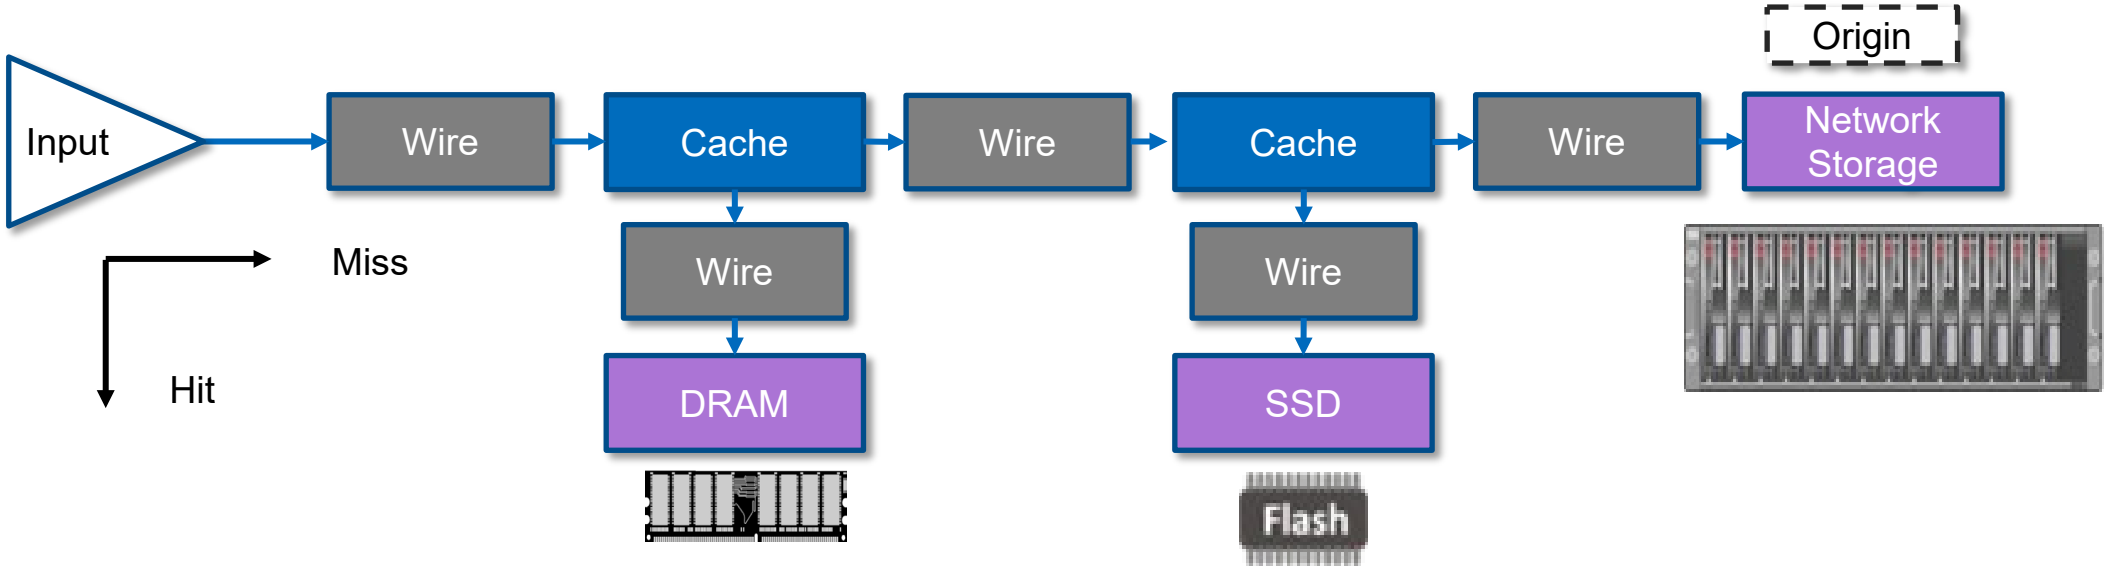

Wire component

- Memory bus, disk controller, network
- Can multiplex and change form of IO request
- Even type of wire can be variable
  - Type of memory bus
  - Hops in network topology
- Delays introduced by wire, contention, queueing

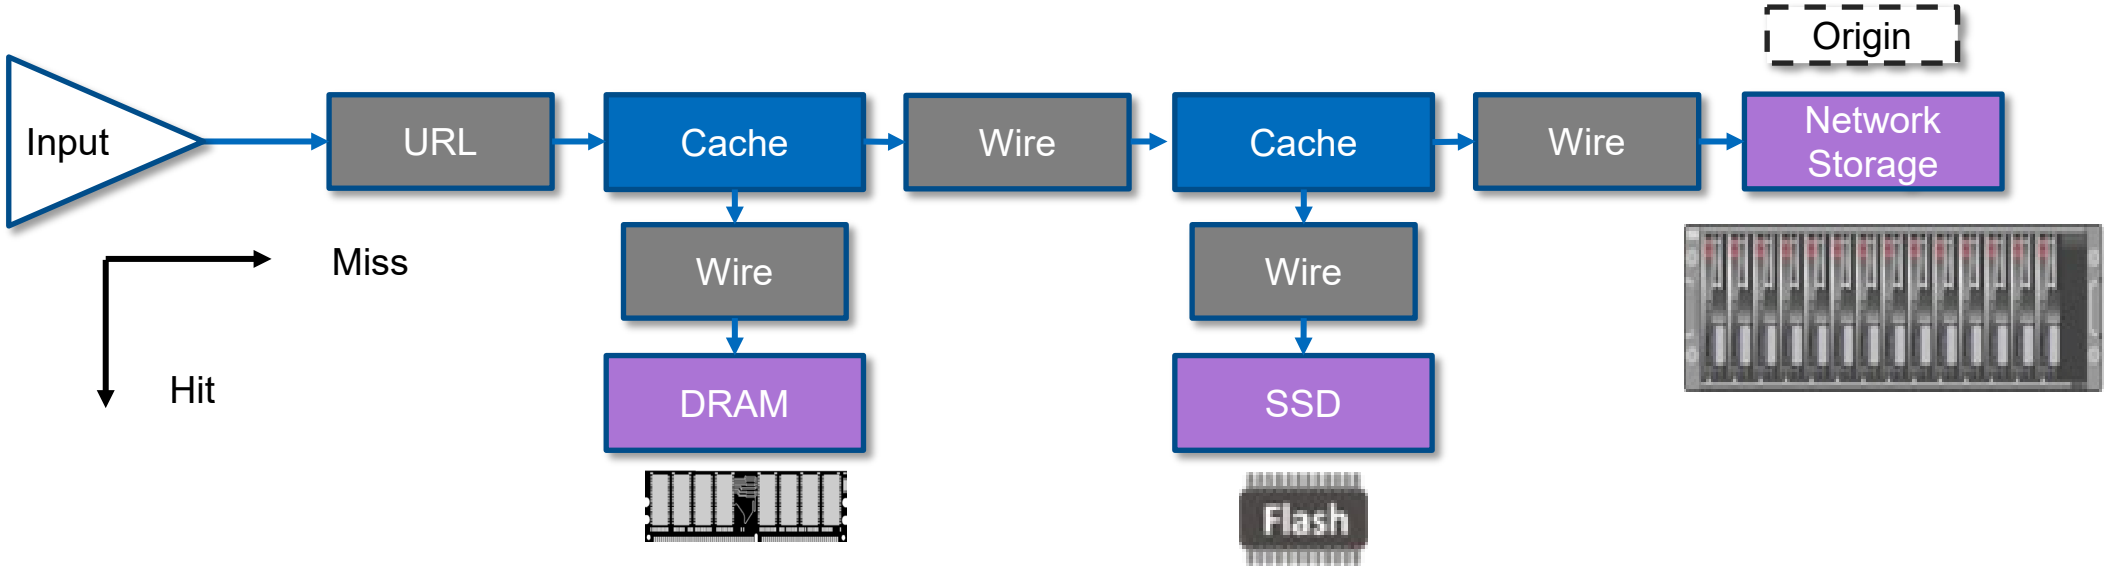

Wire component

- Memory bus, disk controller, network
- Can multiplex and change form of IO request
- Even type of wire can be variable
  - Type of memory bus
  - Hops in network topology
- Delays introduced by wire, contention, queueing

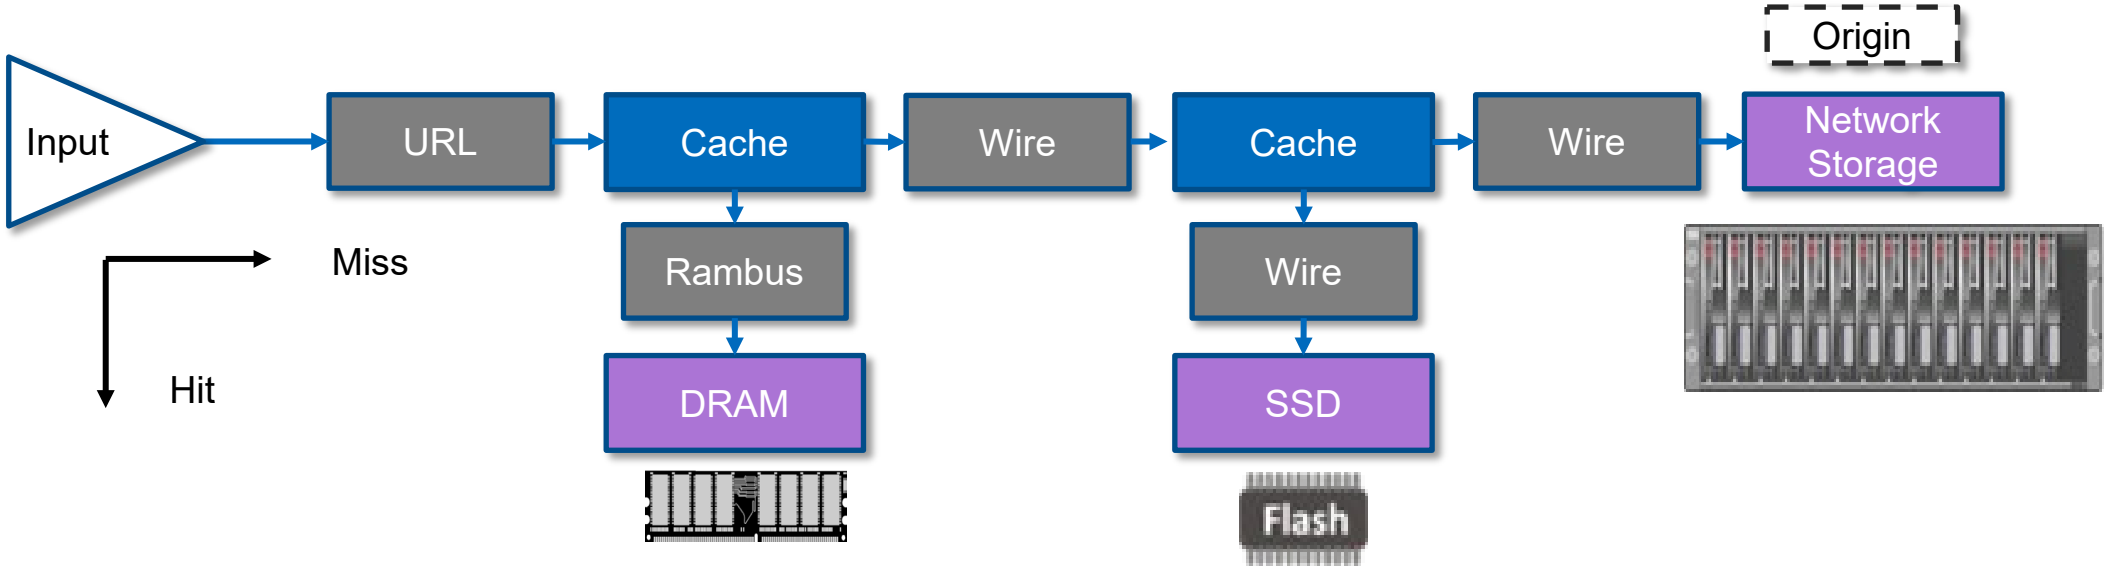

Wire component

- Memory bus, disk controller, network
- Can multiplex and change form of IO request
- Even type of wire can be variable
  - Type of memory bus
  - Hops in network topology
- Delays introduced by wire, contention, queueing

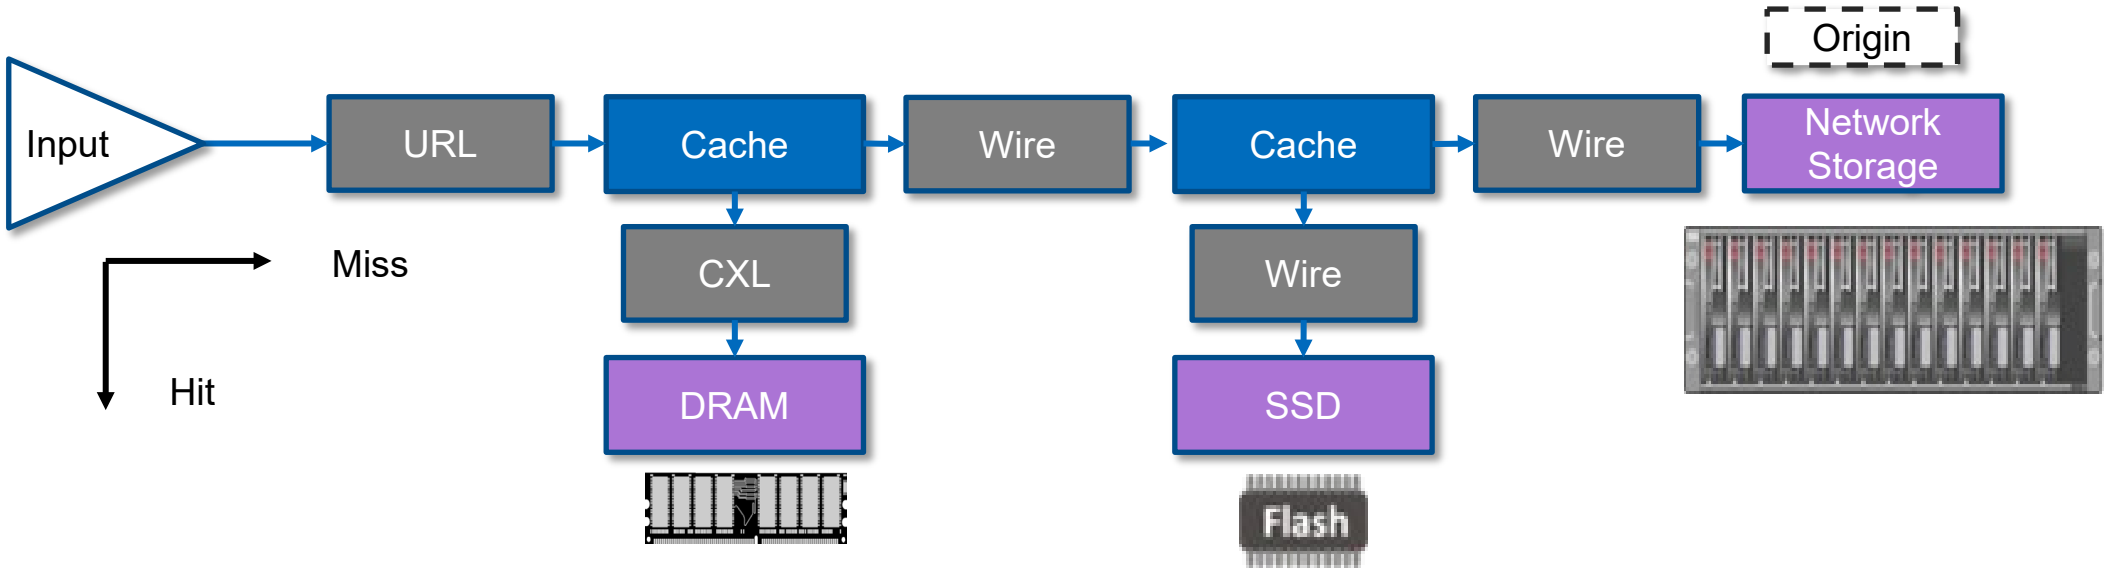

Wire component

- Memory bus, disk controller, network
- Can multiplex and change form of IO request
- Even type of wire can be variable
  - Type of memory bus
  - Hops in network topology
- Delays introduced by wire, contention, queueing

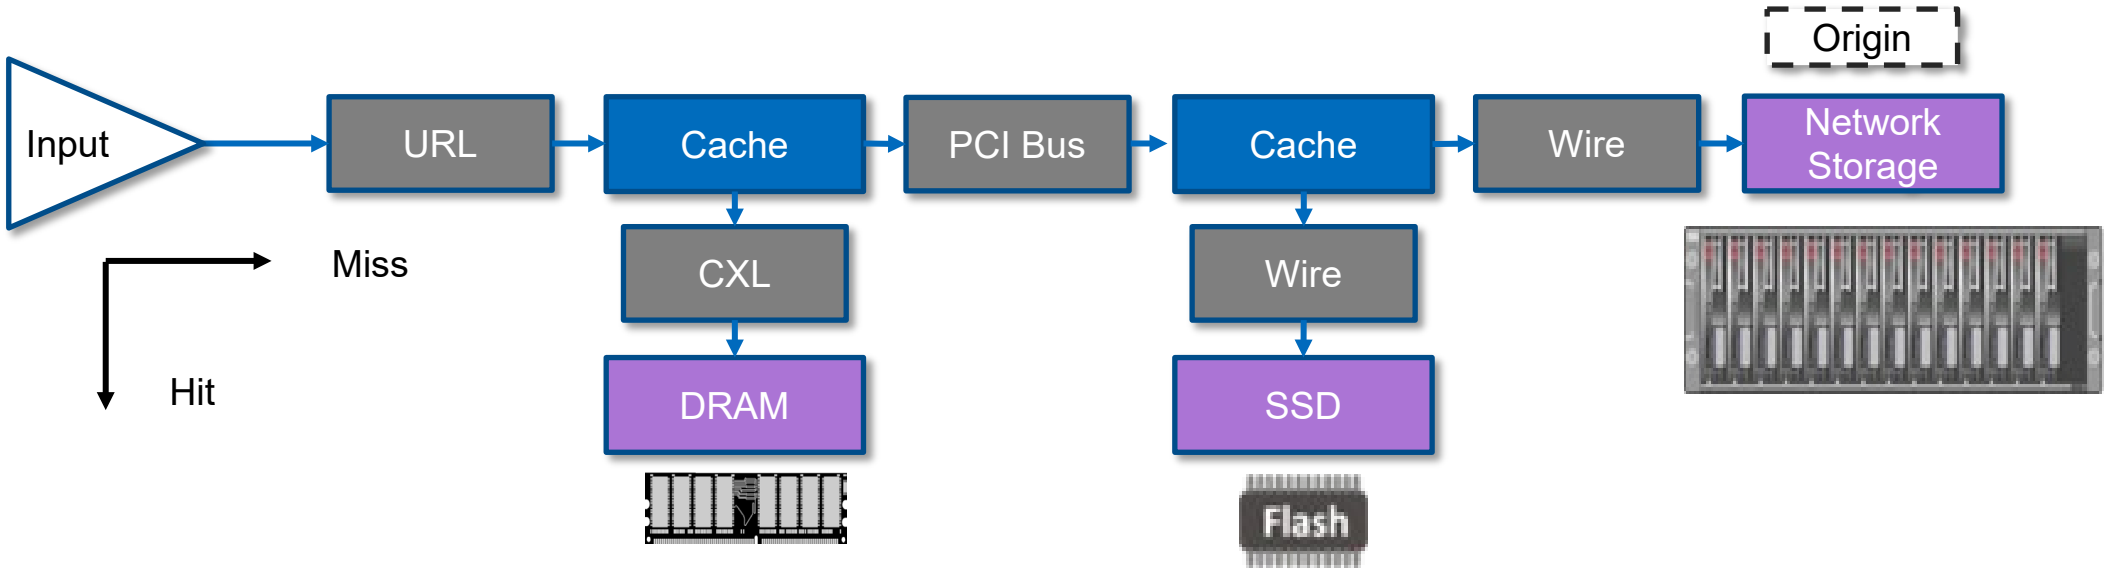

Wire component

- Memory bus, disk controller, network
- Can multiplex and change form of IO request
- Even type of wire can be variable
  - Type of memory bus
  - Hops in network topology
- Delays introduced by wire, contention, queueing

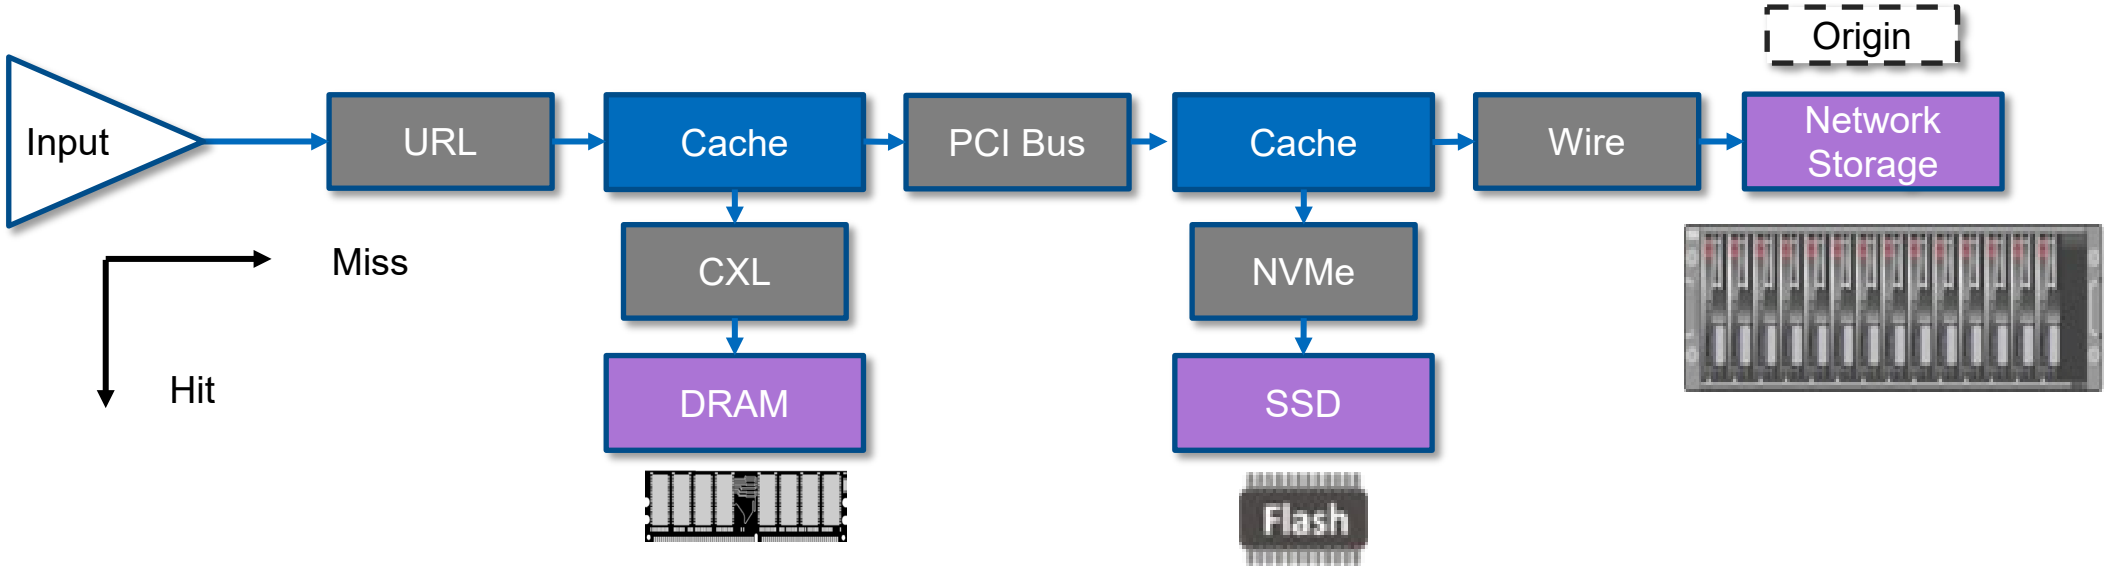

Wire component

- Memory bus, disk controller, network
- Can multiplex and change form of IO request
- Even type of wire can be variable
  - Type of memory bus
  - Hops in network topology
- Delays introduced by wire, contention, queueing

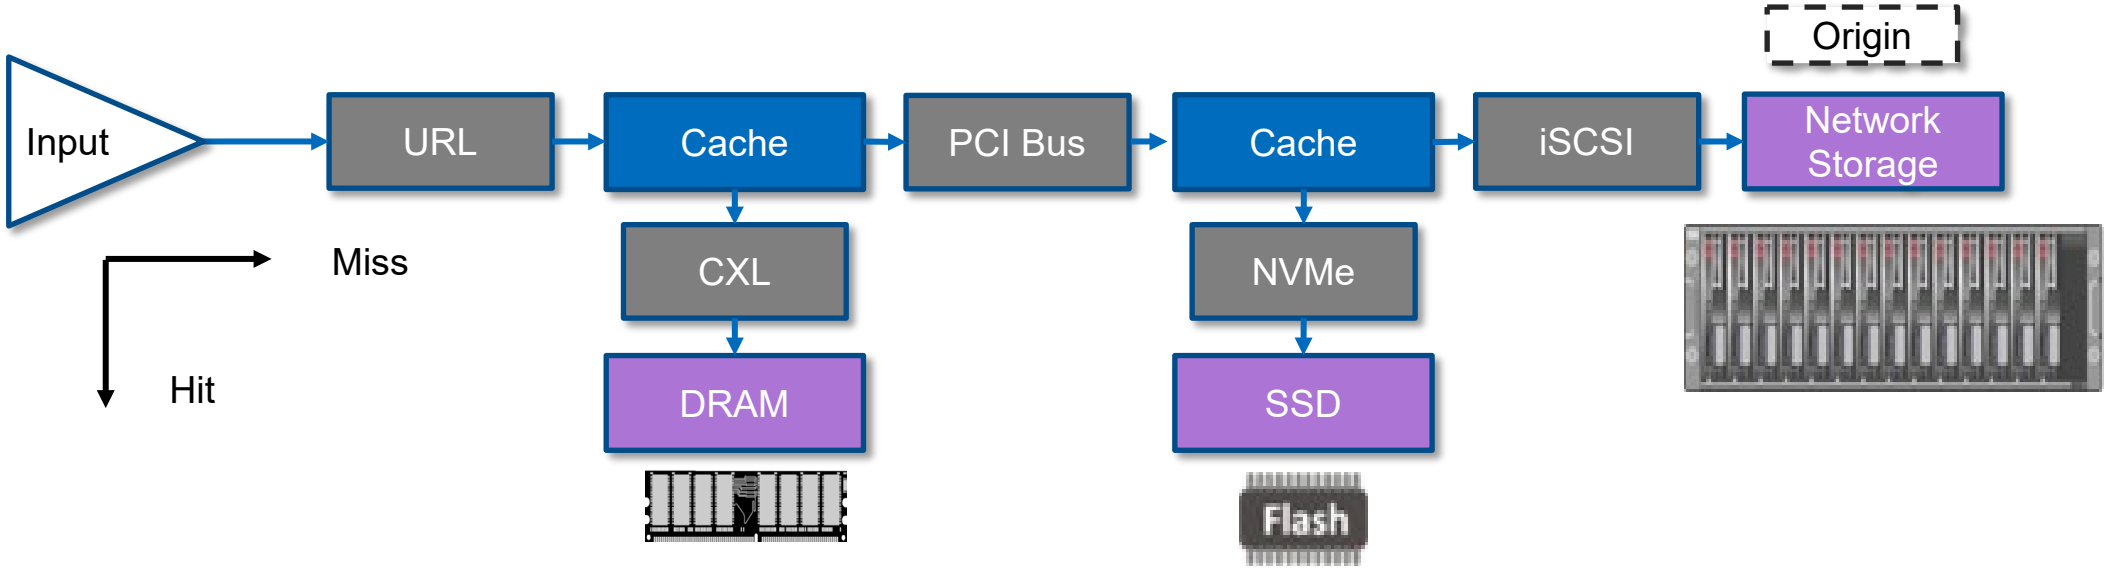

What this means

- **Real performance results**
- **Magnition gets performance success for our client**

Magnition modeling

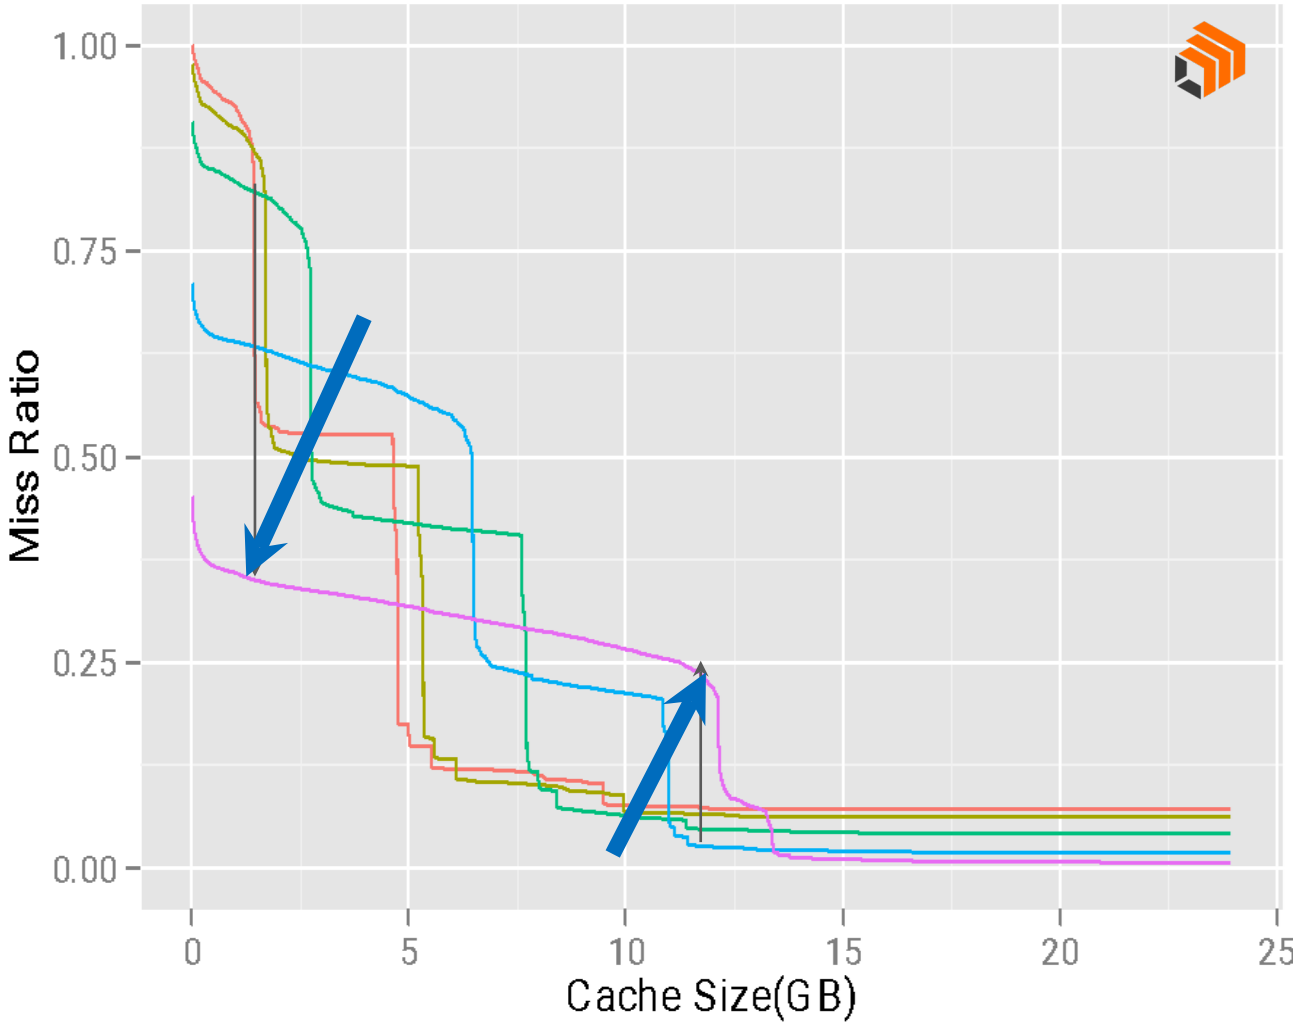

Single workload with Magnition prediction of performance under different policies.

Numerous simulations with variable cache sizes. No hardware required.

Success story

- A company participating in this symposium
- Worked with Magnition over numerous months
- Demonstrated significant performance improvements
- Work integrated into a product they delivered

RESULTS WITH MAGNITION

PROVEN IN MARKET TODAY

As an example, a current customer has achieved the following measurable outcomes with Magnition:

Experiments **per day per engineer**:

- Without Magnition**: 2**
- With Magnition: **50,000+**

Parameter variations tested **before prod release**:

- Without Magnition**: 50**
- With Magnition: **1,000,000+**

Workload performance improvement using our products to find **optimal out-of-the-box settings**: **10-50%+**

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Please take a moment to rate this session.

Your feedback is important to us.

Post-Summit, visit www.snia.org/cms-summit for additional content.

