



# Developing innovative FLASH platforms for the in-memory computing era

Sharad Mehrotra, Ph.D., CEO

25 June 2015

[www.saratoga-speed.com](http://www.saratoga-speed.com)

***Extreme capacity all-flash storage, in-memory performance***

# Developing innovative FLASH platforms for the in-memory computing era

## ABSTRACT

Saratoga Speed is a Silicon Valley startup incubated at Sanmina Corporation. The company is addressing data storage and processing challenges in working with “Big and Fast” data sets at sizes and speeds that are rapidly becoming the norm, but present impossible challenges in today’s computing environments. The initial products from the company offer the world's fastest, densest, and highest capacity all-flash storage arrays. Designed to ingest, process, and store hundreds of terabytes of data at tens of gigabytes per second transfer speeds, the arrays have very low space, power and cooling requirements.

There are two common platform design approaches prevalent in the storage or Big Data infrastructure space today. The first is to select inexpensive and lowest-common-denominator commodity hardware, avoid all custom hardware or firmware, and build a platform exclusively with software. The result is performance, capacity, and density all capped at lowest-common-denominator levels with admittedly more flexible integration and ease-of-use. The second approach favors aggressive use of custom hardware with a thin layer of software for access. While higher levels of performance, capacity and density can be realized with this approach, development cycles are longer.

Saratoga Speed platforms start with commodity hardware and off-the-shelf operating systems. Aggressive design, implementation and integration of proprietary hardware, firmware and software in carefully chosen areas results in capabilities that are far beyond other currently available products.

The talk will describe the challenges in building out a line of products with this design and development approach. Flash packaging, RAID, caching, data path acceleration, service layering and offloads, FPGA integration, and application hosting, are some of the specific areas that will be covered. What worked well, what did not and how we had to adjust the strategy to technical and business reality as we went along.

# Developing innovative FLASH platforms for the in-memory computing era

## SPEAKER BIO

Sharad Mehrotra is the CEO of Saratoga Speed. He has over 25 years of technology and business management experience, including the CEO position in his two previous venture-backed start-ups (Procket Networks, Fabric7 Systems), and executive roles at Hewlett Packard and Sanmina Corporation.

Meeting will be conducted on, Thursday June 25, 2015, 11:30 CDT

at the

Unique Digital Inc. Conference Center  
10595 West Office Drive, Houston, TX 77042

# Company

## Saratoga Speed

- Founded October 2012
- Incubated at Sanmina (NSDAQ:SANM)
- Operating as SANM subsidiary
- First product brought to market in 5-6 quarters

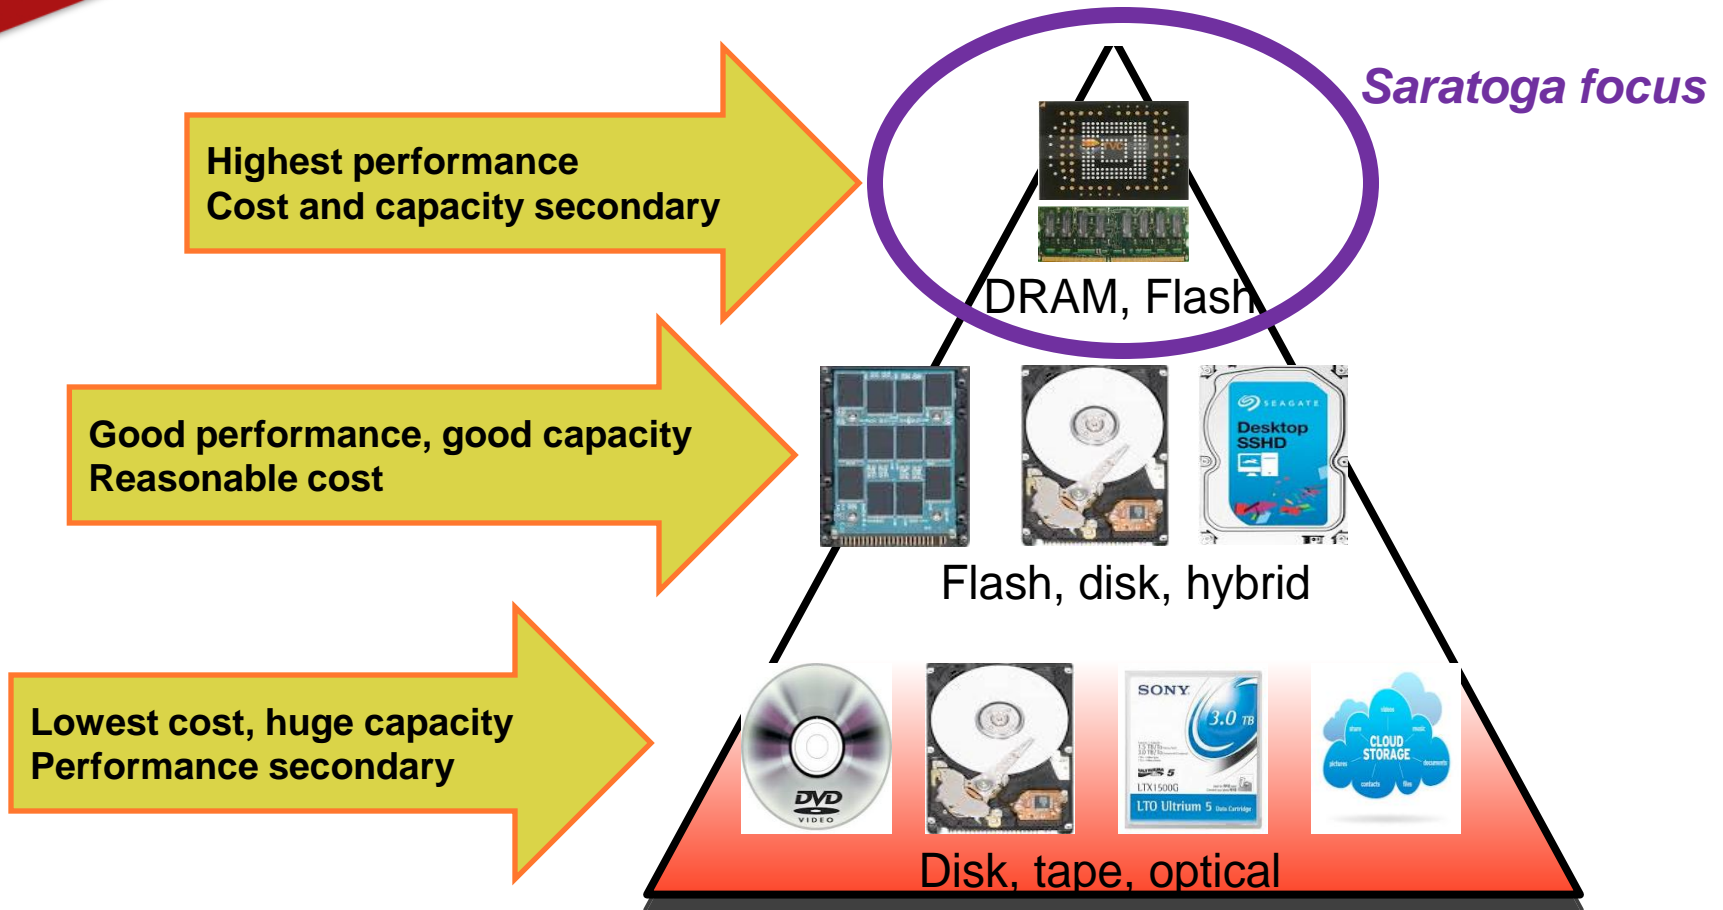
## Sanmina Partnership

- Design services
- Manufacturing
- Facilities and support services: Accounting, Legal, HR and Administrative
- Ready access to working capital to scale revenue production

# Mission

**Eliminate barriers to storing, managing,  
and analyzing big and fast data sets**

# The evolving storage landscape: Performance / Cost / Capacity pyramid



# Big and fast data: Lucrative but hard to analyze

Web activity data (click streams, ad impressions, messaging)

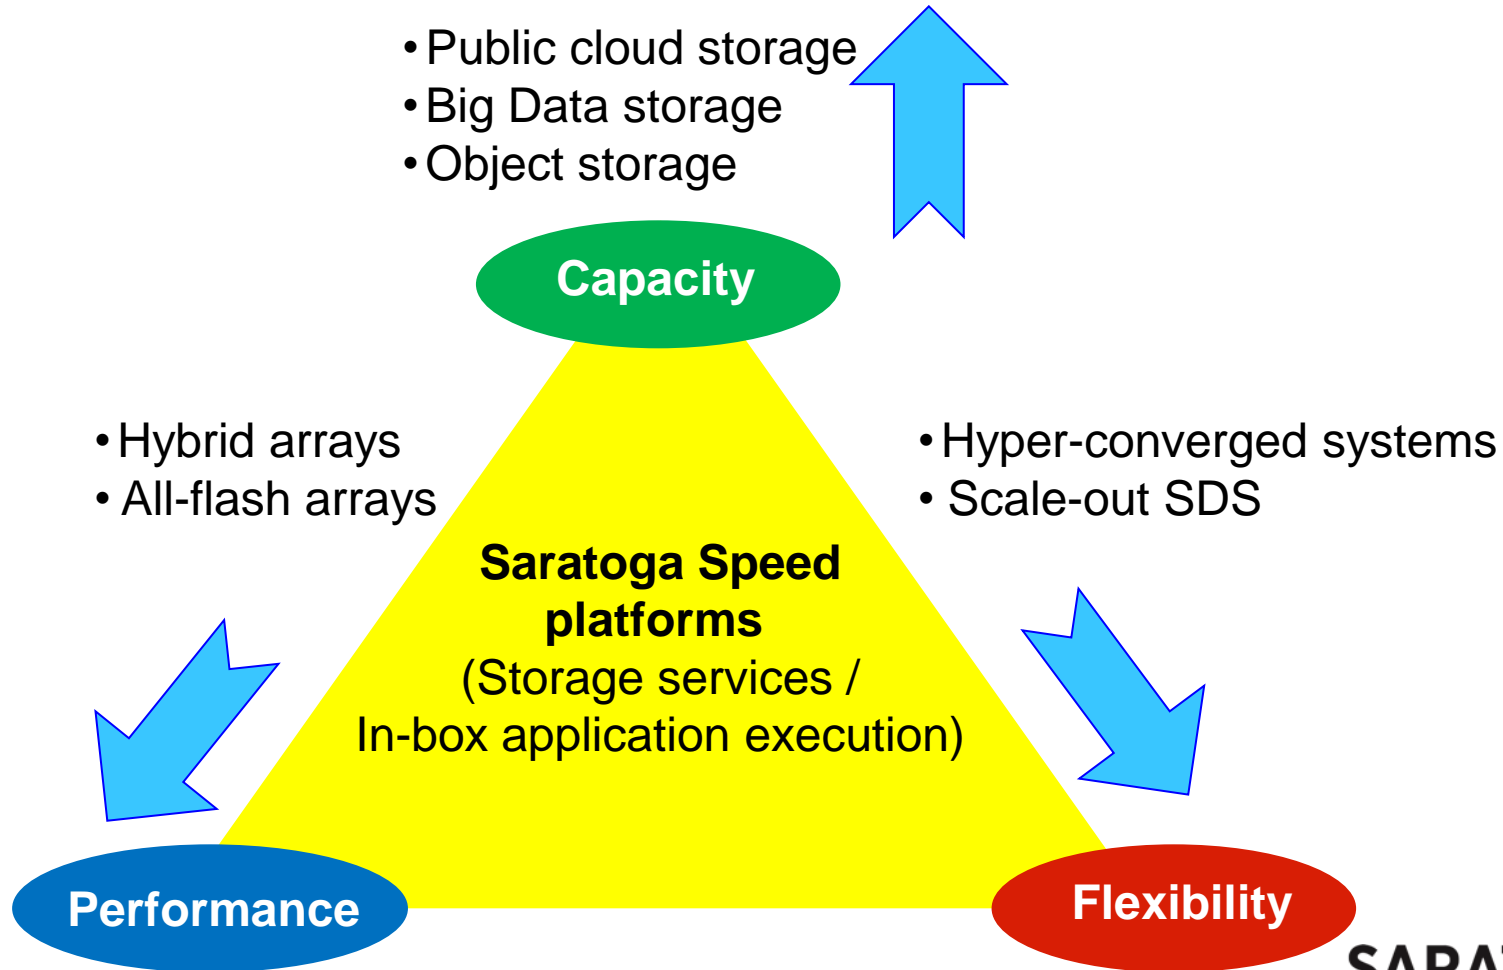
Real-time tick data analysis, compliance, risk and trade models

Telecom data records and content (calls and messaging) analysis

Network traffic inspection and cyber security

- Timely analysis is key
  - Value of the analysis erodes rapidly
  - Large volumes of data discarded
- Hard disk-optimized storage cannot handle the deluge

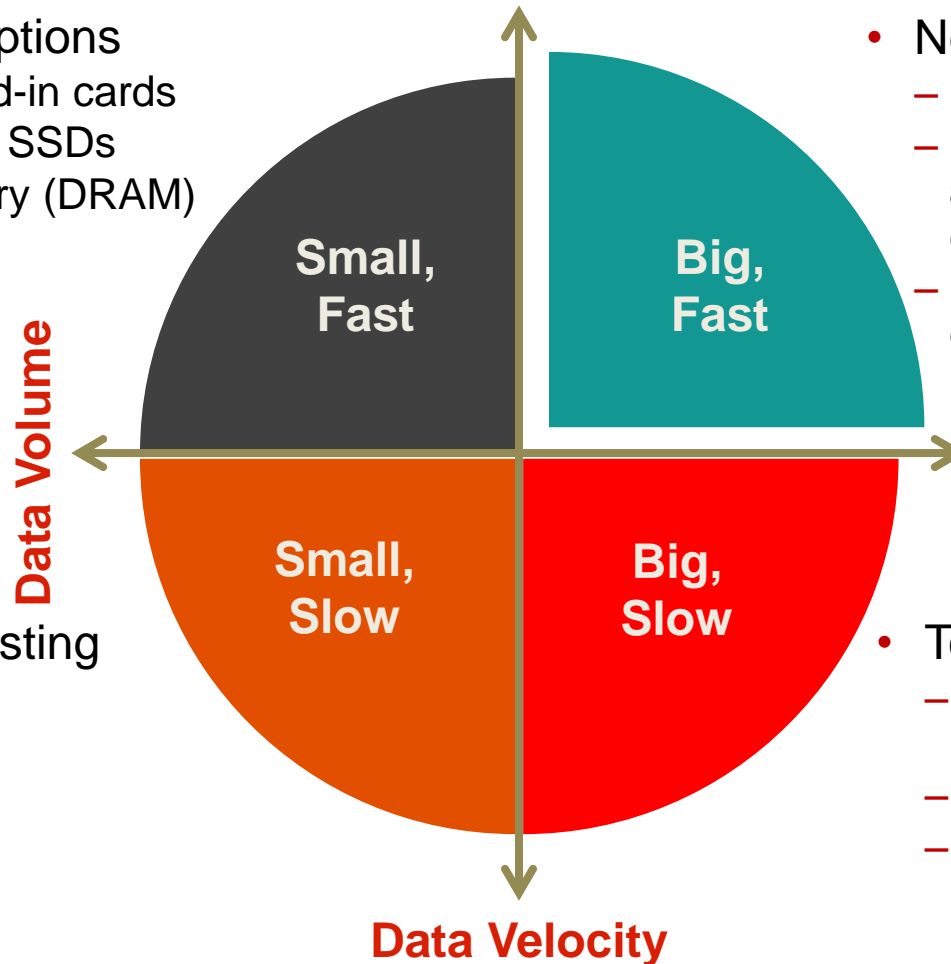
# Saratoga Speed: Nailing the trifecta for Big and Fast data





# Current approaches don't work

- Several options
  - Flash add-in cards
  - In-server SSDs
  - In-memory (DRAM)



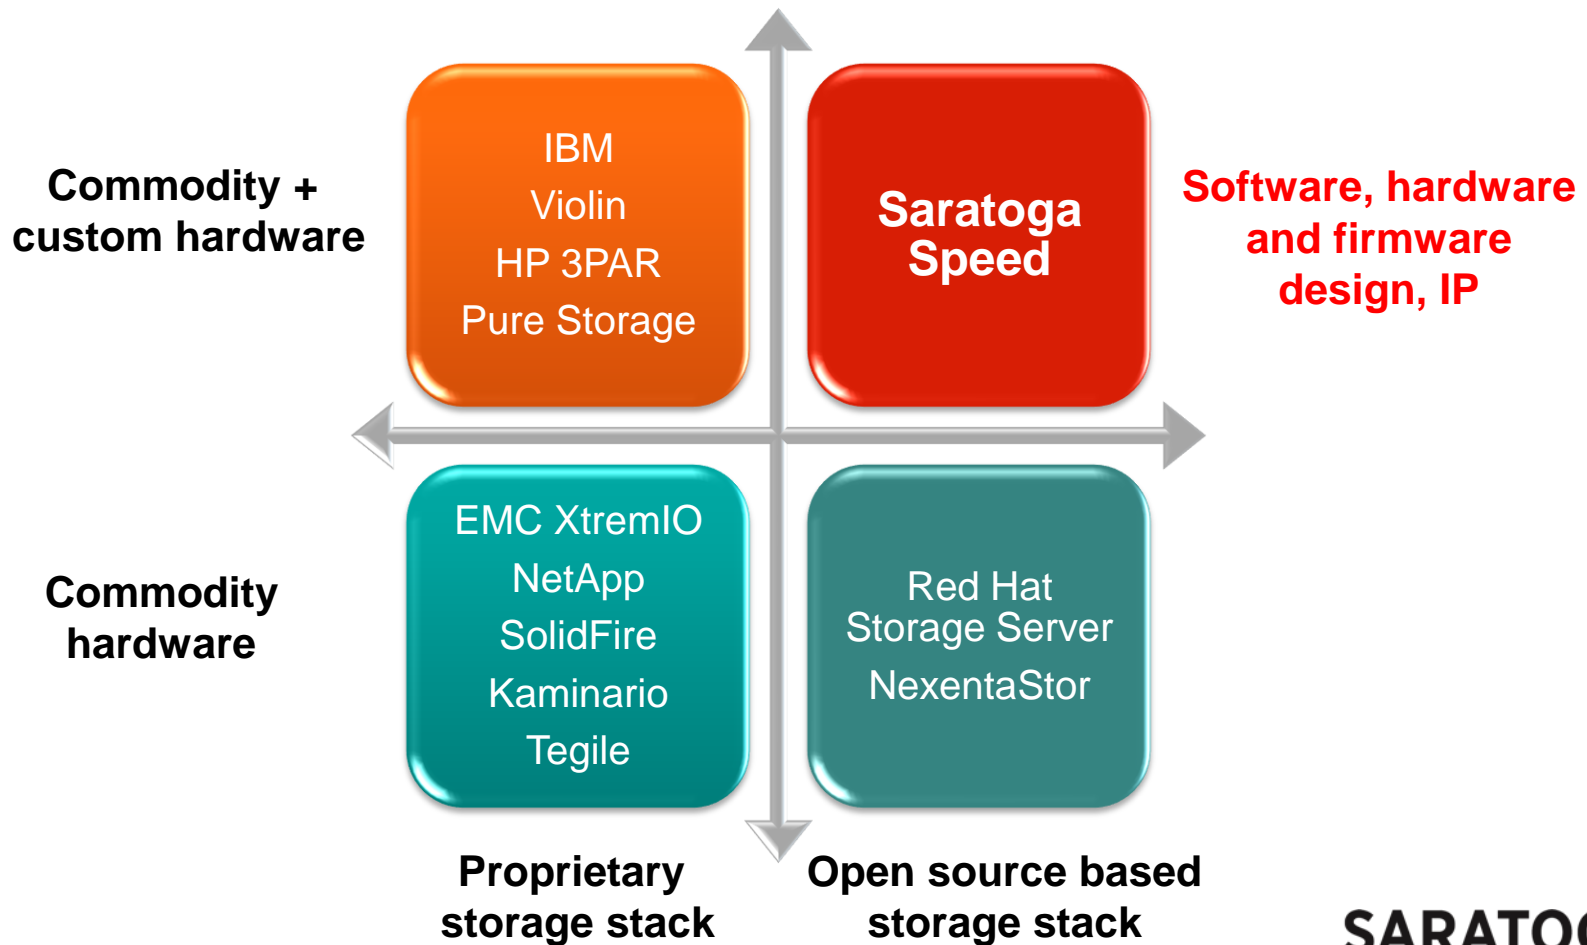
- New solutions required
  - Disks too slow
  - Flash add-in cards and SSDs don't have enough capacity
  - In-memory not economic

- Not interesting

- Today's common case
  - Adequately handled by disk-based systems
  - Commodity scale-out
  - Cloud VM-based infrastructure

# Saratoga Speed: A unique approach to all-flash systems

For advantage in time to market, features, and performance



# First product available, second and third imminent: **RHEL** and **Windows** stacks



## **Altamont XP**

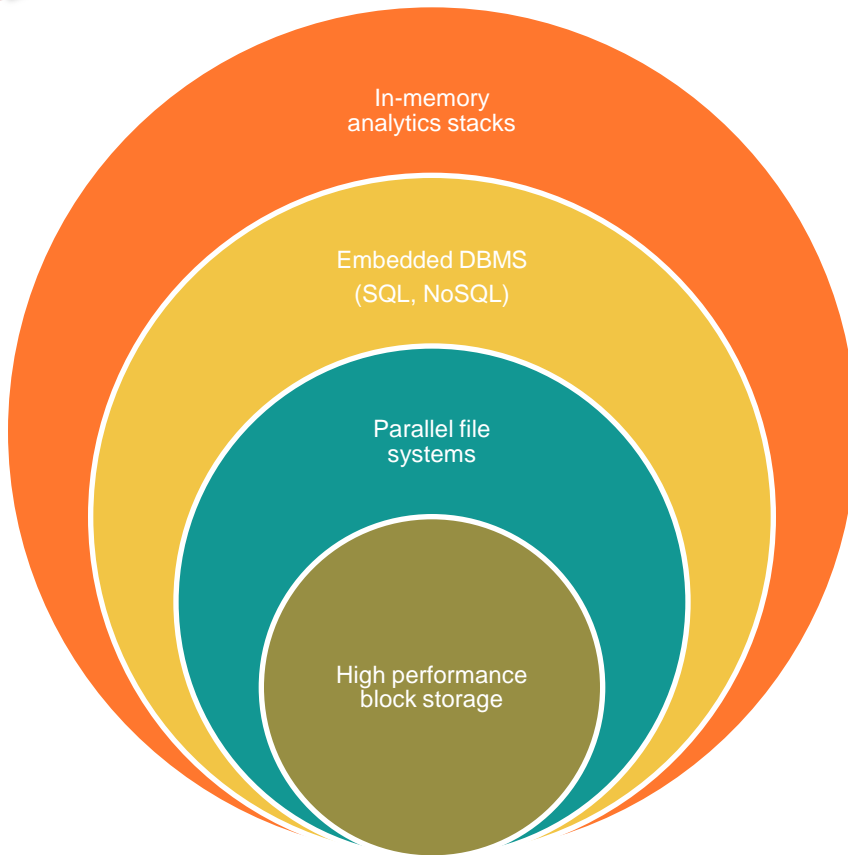
- GA, available now
- Single controller, RHEL
- Tens of TB of usable flash
- 40 GbE, 56 GbIB, 16 GbFC
- RDMA, GPFS, Lustre



## **Altamont XD** **Cascade XP**

- CQ3'15 early access
- Dual controller, RHEL
- Hundreds of TB of usable flash
- 40 GbE, 56 GbIB, 16 GbFC
- RDMA, GPFS, Lustre
- XF-1 Flash Engine

# Key use cases: Systematic expansion

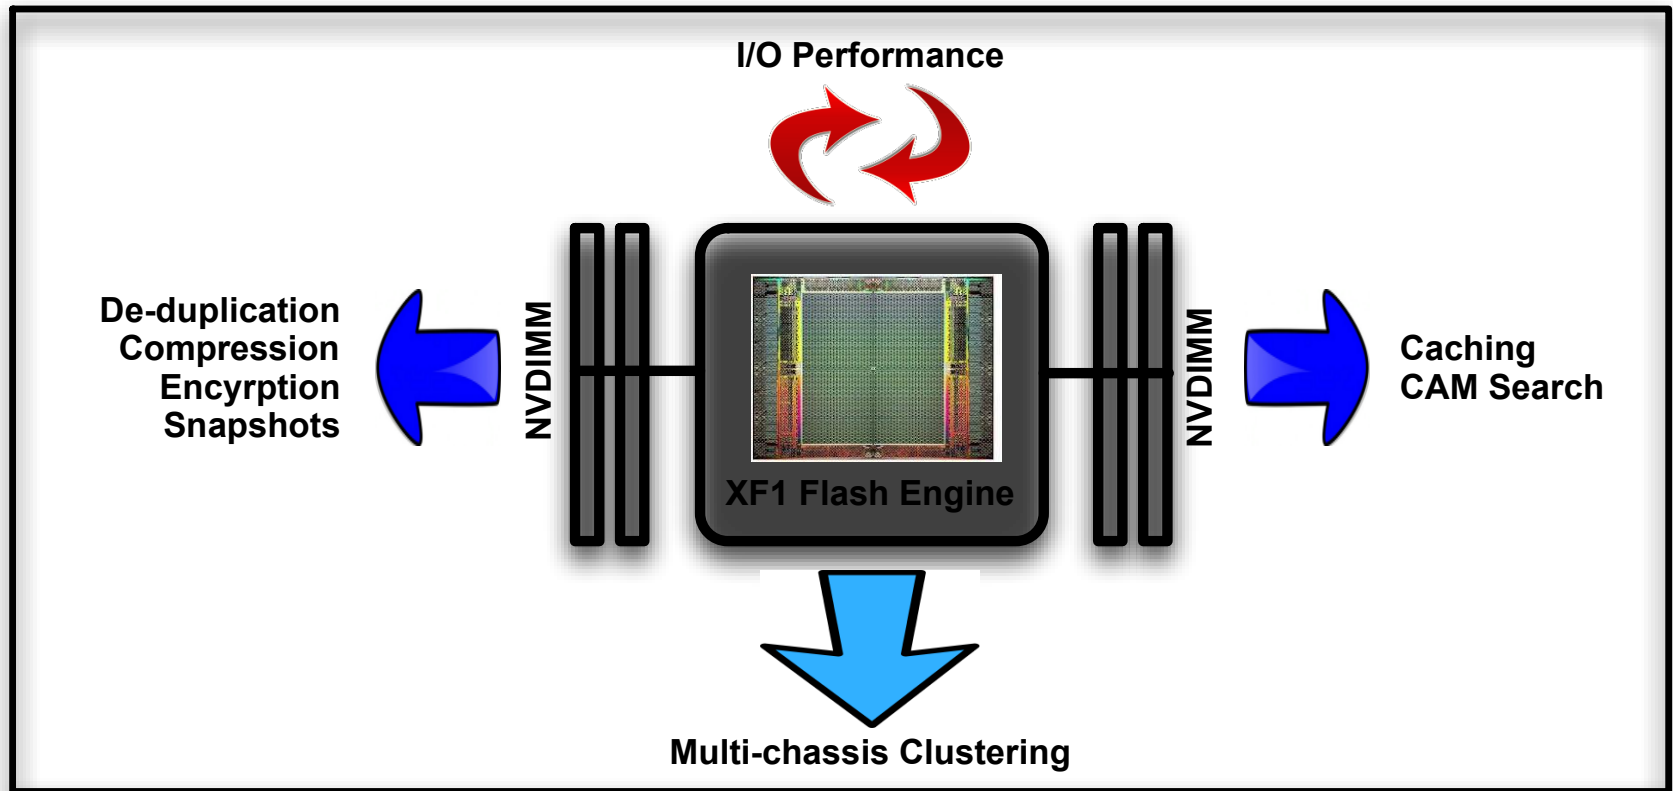


- **High performance block storage** - Online Transaction Processing (OLTP) / Online Analytics Processing (OLAP)
  - Oracle
  - SQL Server
  - SAP ASE
  - SAP HANA / IQ
- **Parallel file systems**
  - IBM Elastic Storage (GPFS)
  - Lustre
- High-performance **storage for virtual machine** or **container** farms
  - VMWare vSphere, Microsoft Hyper-V
  - Docker container clusters
- **Embedded DBMS** all-flash appliances
  - Oracle, SQL Server, MySQL, MongoDB, Cassandra

# Design and implementation challenges

- Flash packaging
- HW RAID
- Caching
- Data path acceleration
- Service layering and Offloads
- RHEL-based software releases
- Active-active controllers with RHEL clustering
- Flash, platform, and system management
- Cloud-based and in-system call-home / data analytics
- Application hosting

# XF-1 Flash Engine: Optimizing performance and flash life



# XF-1 Flash Engine: Optimizing performance and flash life

## Feature

## Benefit

- NVDIMM caching
- CAM search engine
- Write coalescing
- Write linearization
- Offloading and service acceleration
- Traffic management
- Optical interfaces
- FPGA implementation

High-bandwidth, low-latency I/O,  
data protection

Flash life maximization

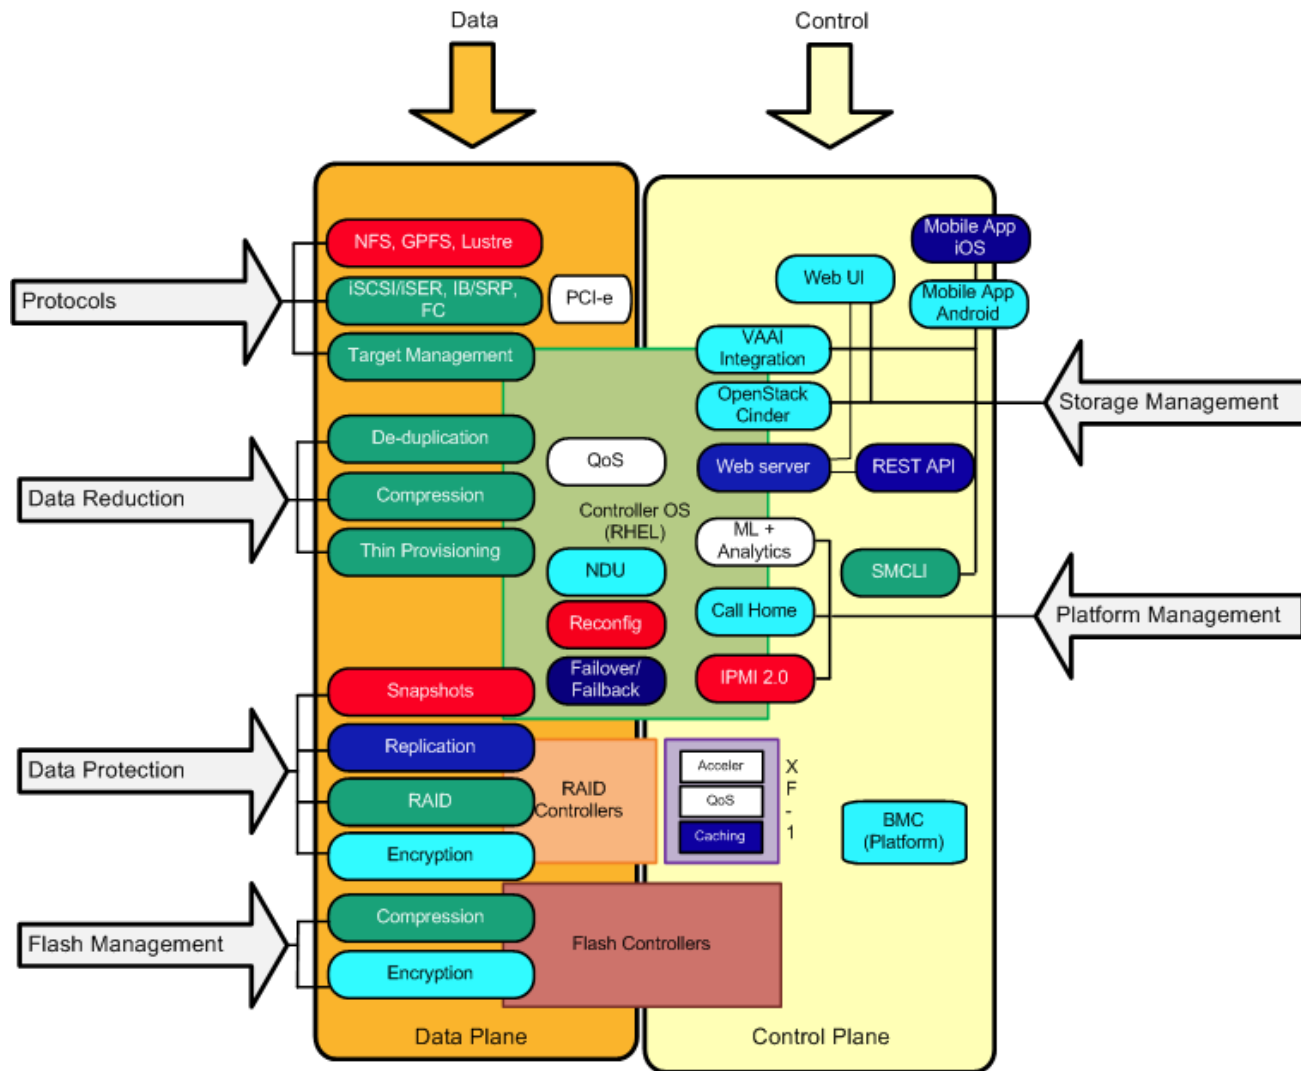
Performance enhancement

QoS guarantees, fair share usage

Built-in interconnects for  
multi-chassis clustering

Investment protection

# RHEL Releases: An iterative approach





# Novel cloud-based call home and machine-learning assisted data optimization

- Machine learning technology to optimize flash lifetime usage and data placement optimization
  - Apply well-understood techniques for online advertising to flash management and data placement
  - Leverage SRI's machine learning IP
  - Gather and transmit flash usage and per-platform system data to cloud back-end
  - Run deep analyses in cloud back-end
  - Acquire classification hints that can be cached on FPGAs
- Cloud back-end in AWS or Azure
  - Leverages work done by Sanmina shop floor data collector team

A large red geometric graphic on the left side of the slide, consisting of several overlapping triangles and polygons that form a stylized arrow or chevron shape pointing to the right.

**Questions?**

[www.saratoga-speed.com](http://www.saratoga-speed.com)

# The evolving storage landscape: Alternatives to traditional Filers / SANs

- Public cloud storage
  - AWS, Azure, Google
- Big Data storage
  - Hadoop distros from Cloudera, HortonWorks, MapR; Red Hat Ceph / Gluster
- Object storage
  - Cleversafe, Scality, Caringo, Amplidata, NexentaEdge
- Scale-out software defined storage (SDS)
  - IBM GPFS, Lustre, EMC ScaleIO, VMware VSAN, DataCore, NexentaStor, etc.
- Hyper-converged systems
  - Nutanix, Simplivity, Gridstore, Springpath, Maxta, Scale Computing, EVO:RAIL solns.
- Hybrid arrays
  - Nimble, Tintri, Tegile, NexGen, Data Gravity, Infinidat, incumbent vendor solutions
- All-flash arrays
  - EMC XtremIO / DSSD, IBM, HP StorServ, Pure Storage, Violin, SolidFire, Kaminario, Coho Data, etc.

**SARATOGA**  
**SPEED**