

# **The Redesigned BaBar Event Store - Believe the Hype**

Adeyemi Adesanya, Jacek Becla,  
Daniel Wang  
(SLAC)

BaBar Computing Group

# BaBar Event Store

- Purpose
  - To provide persistence for billions of physics events across millions of collections for later analysis
- Key requirements
  - Robust and reliable service (close to 100% availability)
  - Abstraction layer to hide persistence implementation
- The current system
  - In successful production since May 1999
  - The largest known database system

# BaBar Event Store: May 1999 - Present

- Highly flexible system
  - Built with little knowledge of use—
    - First of its kind!
  - Generous flexibility has accommodated ongoing adaptation to unforeseen needs
    - Grouping event components using headers
  - Flexibility comes at a price
    - Size matters... especially at this scale!

# The Redesign

- Approach
  - Make use of production experience to reduce size
- Simple techniques for dramatic results
  - Eliminate redundant data by sharing
  - Eliminate obsolete data altogether
  - Reorganize data into more efficient structures
- Side benefits
  - Reduce I/O load → better performance
  - Increase data safety
- By doing this, we also get:
  - Comprehensive code audit (correctness, use cases)
  - New techniques for the analysis model

# Redesign Details(1)

- Reduce overhead: less persistent objects  
(a persistent object has 14 bytes overhead)
  - Combine small persistent objects
  - Pack many strings into one string
- Remove redundant data by sharing
  - Groups of events may have identical (common) attributes
  - Store common values in a single object and share it among many events
- Remove obsolete data

# Redesign Details(2)

- Constraint: minimal impact to end user
  - Backwards compatibility is a must
  - Redesigned schema must be hidden
  - No job reconfiguration should be necessary
  - New system must work with old data

What we want users to say: “It seems a lot faster, but everything works the same”

What we want administrators to say: “Cool! More disk space for logs.”

# Estimated Impact

- Overall size reduction
  - From 2.2kB to 0.5kB per event (nav.)
  - From 45kB to 5.6kB per tag descriptor
  - From 7kB to 4kB per collection
- Performance speedups
  - Significant, but tough to quantify
  - Requests are fewer and for less data

# Summary

- Redesign builds on the current, proven system
- Simple techniques produce significant gains
- Soon to be released