



Education

RETAINING INFORMATION FOR 100 YEARS

Mary Baker, HP Labs
Roger Cummings, Symantec Research Labs

- The material contained in this tutorial is copyrighted by the SNIA.
 - Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
 - This presentation is a project of the SNIA Education Committee.
 - Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
 - The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.
- NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.**

➤ Retaining Information for 100 years

- ◆ Many organizations now have a requirement to preserve large volumes of digital content indefinitely into the future, and to maintain access for reasons such as medical treatment decisions, retention of intellectual property, and appreciation of cultural and scientific history. Frequent news stories cover organizations' failures to be able to do this, such as the near loss of original video/data of the first Moon landing, eventually recovered from a set of 14-inch tape reels found in a dusty Australian basement.
- ◆ This session will focus on the most important questions in long-term digital preservation and will demonstrate why it is still so difficult. We will propose how the storage industry can help its customers preserve and use their digital content over the lifetimes that they expect from past experience with physical and analog assets, lifetimes that can greatly exceed those of any single digital storage device or storage technology.

Outline

➤ Introduction

- ◆ Why we need digital preservation
- ◆ The goals of digital preservation

➤ Threats

- ◆ Threats to long-term digital content
- ◆ How long-term and short-term threats differ
- ◆ Why it is hard to address these threats

➤ Current status

- ◆ Best practices
- ◆ A current project
- ◆ Open problems and opportunities
- ◆ Please give us feedback

Why we need digital preservation

- Regulatory compliance and legal issues
 - ◆ Sarbanes-Oxley, HIPAA, FRCP, intellectual property litigation
 - Emerging web services and applications
 - ◆ Email, photo sharing, web site archives, social networks, blogs
 - Many other fixed-content repositories
 - ◆ Scientific data, intelligence, libraries, movies, music
-

- Responses to 100 Year Archive Requirements Survey
 - ◆ 68% of organizations had requirements for over 100 years
 - ◆ 83% of organizations had requirements for over 50 years



Goals of digital preservation

- Digital assets stored now should remain
 - ◆ accessible
 - ◆ usable
 - ◆ undamaged

- for as long as desired – beyond the lifetime of
 - ◆ any particular storage system
 - ◆ any particular storage technology

- and at an *affordable cost*

Outline

➤ Introduction

- ◆ Why we need digital preservation
- ◆ The goals of digital preservation

➤ Threats

- ◆ Threats to long-term digital content
- ◆ How long-term and short-term threats differ
- ◆ Why it is hard to address these threats

➤ Current status

- ◆ Best practices
- ◆ A current project
- ◆ Open problems and opportunities
- ◆ Please give us feedback

Threats to long-term digital assets

- Large-scale disaster
- Human error
- Media faults

- Component faults
- Economic faults
- Attack
- Organizational faults

Long-term content suffers from more threats than short-term content

- Media/hardware obsolescence
- Software/format obsolescence
- Lost context/metadata

Threats to long-term digital assets

- Large-scale disaster
- Human error
- Media faults ←
- Component faults
- Economic faults
- Attack
- Organizational faults



- Media/hardware obsolescence
- Software/format obsolescence
- Lost context/metadata

Threats to long-term digital assets

- Large-scale disaster
- Human error
- Media faults

- Component faults
- Economic faults
- Attack ←
- Organizational faults




- Media/hardware obsolescence
- Software/format obsolescence
- Lost context/metadata

Threats to long-term digital assets

- Large-scale disaster
- Human error
- Media faults

- Component faults
- Economic faults
- Attack
- Organizational faults



- Media/hardware obsolescence ←
- Software/format obsolescence
- Lost context/metadata

Even preserving just the bits is hard

- Large scale & long time periods are a problem
- 1 petabyte, 50 years, 50% probability of no damage
 - ◆ Sounds reasonable, doesn't it?
- That's a bit half-life of 10^{17} years
 - ◆ A million times the age of the universe
 - ◆ Even improbable events will have an effect
- Now try to keep
 - ◆ The bits usable
 - ◆ The information reusable
 - ◆ The applications usable
- Preserve just the bits (physical preservation)?
 - ◆ Can't interpret the content
- Focus only on the logical aspects (logical preservation)?
 - ◆ The bits have been trashed

Outline

➤ Introduction

- ◆ Why we need digital preservation
- ◆ The goals of digital preservation

➤ Threats

- ◆ Threats to long-term digital content
- ◆ How long-term and short-term threats differ
- ◆ Why it is hard to address these threats

➤ **Current status**

- ◆ Best practices
- ◆ A current project
- ◆ Open problems and opportunities
- ◆ Please give us feedback


Key ideas behind best practices

- **Replicate content**
 - ◆ If one copy is damaged, can repair from another
 - ◆ It's not enough to make a single “super reliable” copy
- **Avoid correlated failures**
 - ◆ Not just geographic, but administrative, platform, etc...
 - ◆ Heterogeneity helps avoid correlations
 - ◆ Must balance this with cost and administrative hassle
- **Find & fix (if possible) latent faults before damage grows**
 - ◆ Some faults don't announce themselves
 - ◆ Latent faults can occur at all layers, physical and logical
 - ◆ Must look for silent damage/problems proactively: “audit”
 - ◆ This means the content must be accessible!
- **Use widely understood standards**
 - ◆ Help customers avoid metadata and format traps
 - ◆ Help customers migrate content to (your!) new technologies

Example: audited replicated archive

Baker et al., "A Fresh Look at the Reliability of Long-term Digital Storage." *EuroSys 2006*.

Reliability vs. Auditing



Best practices will vary over time

- We can't predict what will change – we only know it will
 - ◆ Ability to evolve is most important aspect of digital preservation
- This means processes are key
 - ◆ Must ensure our preservation processes are evolvable
 - ◆ Current processes are the first step in an iterative solution
 - > They get us to the next step
 - > At that point we will likely need new processes to take over
 - ◆ Physical preservationists aim for undoable transformations
 - ◆ Widely understood standards make process evolution easier
- A good archive is almost always in motion
 - ◆ Migrating, auditing, re-keying, etc.
 - ◆ **Digital preservation is not a static activity!**
 - ◆ You can't just “do it and be done with it”

Best practices will vary by context

- What do we preserve?
 - ◆ Bits? Applications? Logical connections? Context? Etc.?
- Preserve what the customer in that domain wants
 - ◆ Different domains/industries/organizations need different things
 - > Static versus dynamic content
 - > Self-contained content versus many external dependencies
 - > Different levels of fidelity and context
 - ◆ Example: digital copy of old book
 - > Just copy the words?
 - > Reproduce wear and tear on the paper?
 - > What about the political context in which it was read?
- Can't predict the eventual use of the material
- Affordability may force some decisions

Which methods are best?

➤ Do we use

- ◆ Virtual machines?
- ◆ Emulation?
- ◆ Canonical formats?
- ◆ Self-describing formats?
- ◆ Standardized data models?
- ◆ Loss-tolerant formats?
- ◆ Format migration?
- ◆ Preservation of ancient equipment?

➤ Yes: all could play a role for different domains

- ◆ Some can be very expensive
- ◆ Workable processes vary across organizations/domains

Logical preservation efforts

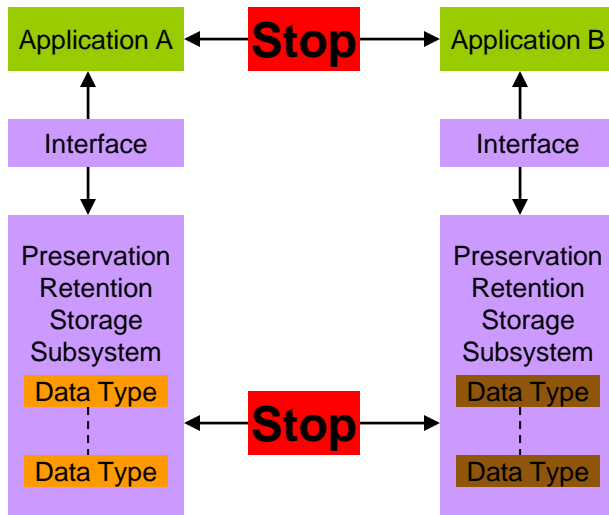
- Long-term Retention Technical Working Group
 - ◆ Both physical and logical preservation efforts
 - ◆ Migration is a potentially affordable approach for both
 - ◆ <http://www.snia.org/apps/org/workgroup/ltrtwg/>
- Logical migration
 - ◆ A means to interpret content into the future
 - ◆ We need tested, affordable, scalable solutions
- Self-contained Information Retention Format
 - ◆ (SIRF)
 - ◆ A CASPAR collaboration (EU co-funded program)
 - ◆ CASPAR = Cultural, Artistic & Scientific knowledge for Preservation, Access & Retrieval

SIRF: logical container format

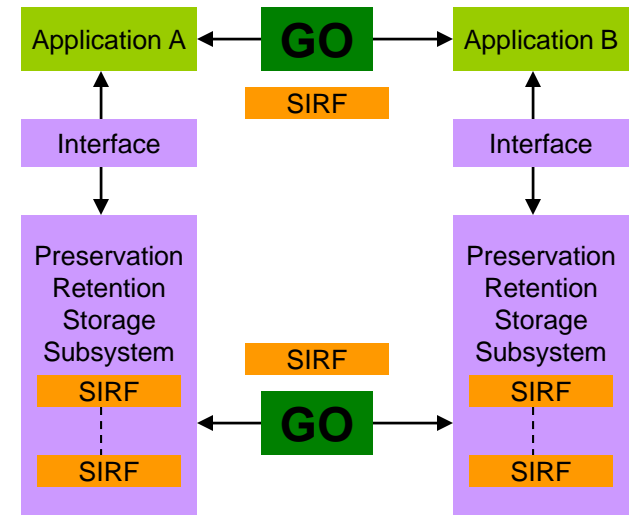
- Appropriate for long-term storage of digital information
- Logical data format of a mountable unit
 - ◆ File system, block device, stream device, object store, tape, etc.
- Includes a cluster of “interpretable” preservation objects
 - ◆ Self-describing – can be interpreted by different systems
 - ◆ Self-contained – all interpretation data contained in object cluster
- Facilitate transparent migration for long-term preservation
 - ◆ Logical
 - ◆ Physical
- Several implementations envisioned
 - ◆ OAIS (Open Archival Information System) ISO standard
 - ◆ XAM (eXtensible Access Method)
 - ◆ Others

Problem SIRF addresses

Without SIRF



With SIRF



Cannot move cluster of preservation objects between systems without help	Can move cluster of preservation objects between systems by itself
Only original application that wrote the preservation objects can read and interpret them	Any SIRF compliant application can read and interpret preservation objects
Need export and import processes	No need for export and import processes
Preservation objects cannot be sustained for long-term	Preservation objects survive longer

Related tutorials to check out



SNIA Tutorial: Cloud Archiving



SNIA Tutorial Track on Green Storage

For more information

- **Survey of data retention problems and requirements**
 - ◆ 100 Year Archive Task Force, SNIA Data Management Forum, “100 Year Archive Requirements Survey,” January 2007.
- **Measurements, modeling of storage failure**
 - ◆ M. Baker, et al., “A Fresh Look at the Reliability of Long-term Digital Storage.” EuroSys’06.
 - ◆ E. Pinheiro et al., “Failure Trends in a Large Disk Drive Population.” Usenix FAST’07.
 - ◆ B. Schroeder et al., “Disk failures in the real world: What does an MTTF of 1,000,000 hours mean too you?” Usenix FAST’07.
 - ◆ L. Bairavasundaram, et al., “An analysis of data corruption in the storage stack.” Usenix FAST’08.
 - ◆ W. Jiang, et al., “Are Disks the Dominant Contributor for Storage Failures? A Comprehensive Study of Storage Subsystem Failure Characteristics.” Usenix FAST’08.
 - ◆ A. Krioukov et al., “Parity Lost and Parity Regained.” Usenix FAST’08.
- **SNIA LTR TWG compiling a knowledge base (KB) for many topics**
 - ◆ Standards such as OAIS
 - ◆ Pointers to work on related problems (authentication, power management, etc.)
 - ◆ Project and deployment descriptions (British Library, CASPAR, LOCKSS)
 - ◆ Intent is to publish KB as a “white paper”

Some open problems & opportunities

- Logical and physical migration
 - ◆ Contribute to the activity in this area!
- Failure data
 - ◆ What are all the ways we really lose content?
- Reliability modeling
 - ◆ Holistic models to reason about probabilities of content loss
- Accelerated aging
 - ◆ How do we know if we've been successful?
- Dealing with secrets for long periods of time
 - ◆ Secrets can be the hardest things to preserve
- Long-term cost modeling
 - ◆ What is the cost to preserve this document for 100 years?
- 3rd-party validation of storage SLAs
 - ◆ Ways to tell that a preservation service is meeting its promises
- Choosing what to preserve
 - ◆ Can/should we save everything? If not, how do we choose?

Summary

- Digital preservation is important now
 - ◆ And is becoming more so
- Best practices center around
 - ◆ Replicating content
 - ◆ Avoiding correlated failures
 - ◆ Auditing for latent damage
 - ◆ Choosing formats/processes that are easy to evolve
- Preservation requires the ability to evolve
 - ◆ Current choices make future evolution harder or easier
- Both logical and bit preservation are important
 - ◆ And remain hard in terms of scalability and affordability
 - ◆ Several interesting projects are underway
- There are many open, critical problems to work on
 - ◆ Please join us!

- This tutorial has been developed, reviewed and approved by members of the SNIA Long Term Retention Technical Working Group (LTR TWG)
 - ◆ **Mission**
 - › The TWG will lead storage industry collaboration with groups concerned with, and develop technologies, models, educational materials and practices related to, data & information retention & preservation.
 - ◆ **Charter**
 - › The TWG will ensure that SNIA plays a full part in addressing the "grand technical challenges" of long term digital information retention & preservation, namely both physical ("bit") and logical preservation.
 - › The TWG will generate reference architectures, create new technical definitions for formats, interfaces and services, and author educational materials. The group will work to ensure that digital information can be efficiently and effectively preserved for many decades, even when devices are constantly replaced, new technologies, applications and formats are introduced, consumers (designated communities) often change, and so on.
- http://www.snia.org/tech_activities/workgroups
- **Please join us!**

- Please send any questions or comments on this presentation to SNIA: trackdatamgmt@snia.org

**Many thanks to the following individuals
for their contributions to this tutorial.**

- SNIA Education Committee

**Mary Baker
Wendy Betts
Simona Cohen
Roger Cummings
Sam Fineberg
Michael Fishman
Annie Foong
Sami Iren
Petros Maniatis**

**Rob Peglar
Michael Peterson
Jeff Porter
Bob Rogers
David Rosenthal
Ramin Samadani
Mehul Shah
Irwin Sobel
Gary Zasman**

Additional material

Can we model long-term reliability?

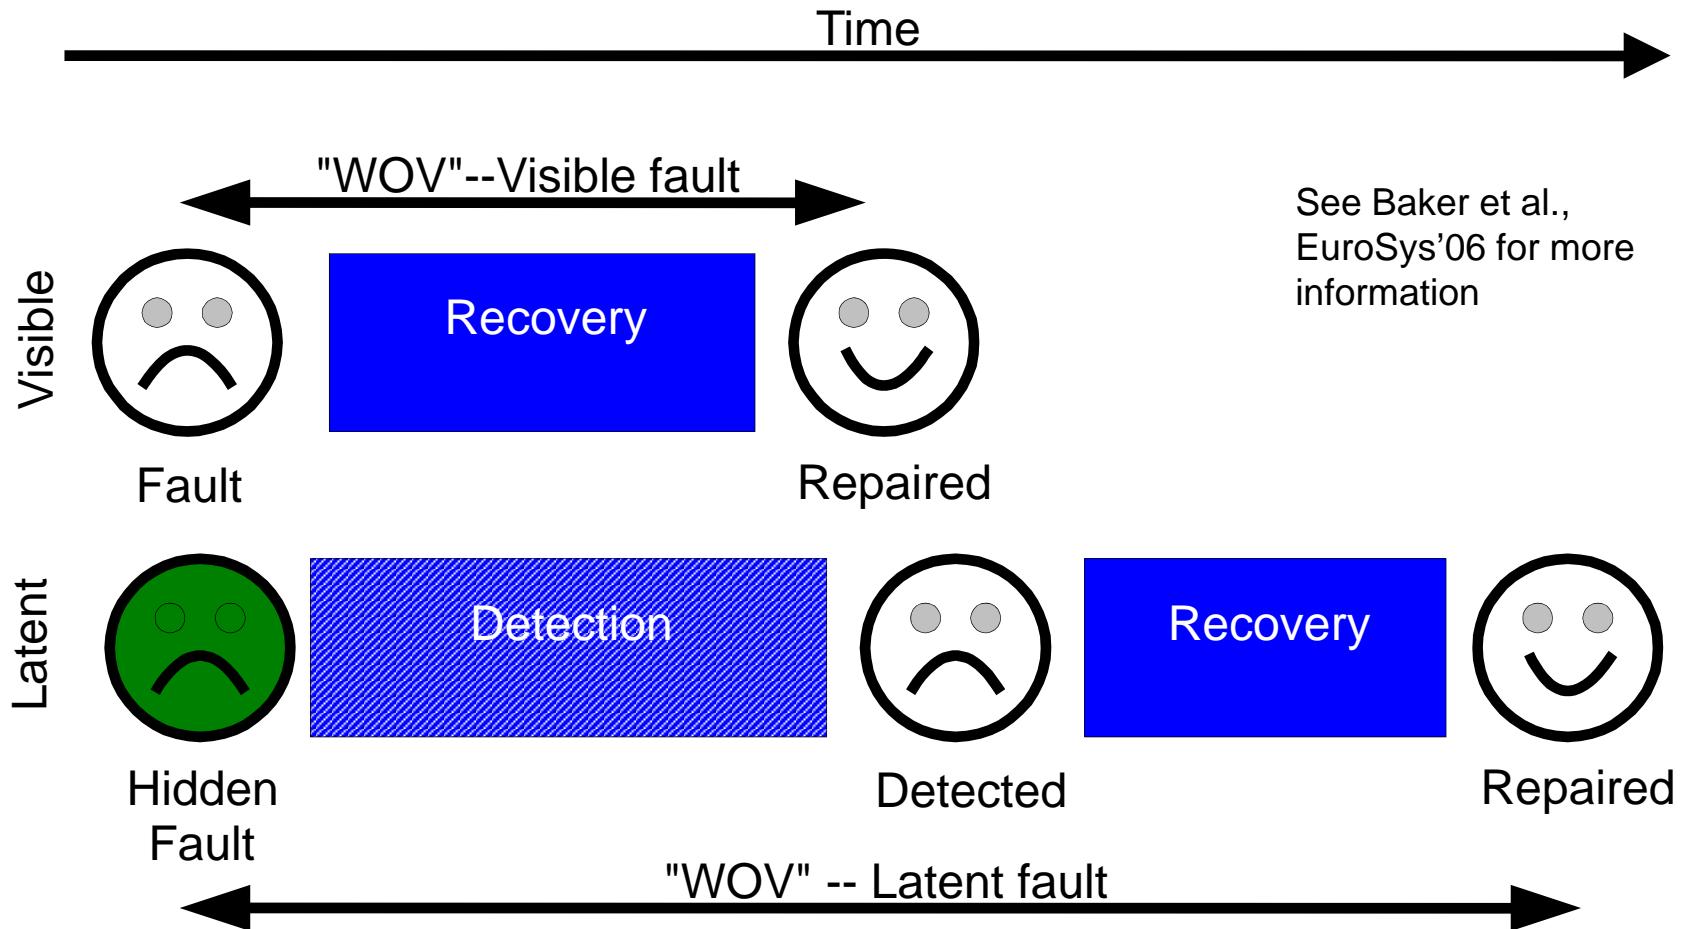
- Abstract reliability model for replicated data
 - ◆ Applies to all units of replication
 - ◆ Applies to many types of faults
- Extend RAID model
 - ◆ Account for latent as well as visible faults
 - ◆ Account for correlated faults: temporal and spatial
- Simple, coarse model
 - ◆ Suggest and compare strategies (choose trade-offs)
 - ◆ Point out areas where we need to gather data
- *Not for exact reliability numbers*

A current approach

- Start with two replicas, then add more
- Derive MTTDL of mirrored data in the face of
 - ◆ Both immediately visible and latent faults
- Mirrored data is unrecoverable
 - ◆ If copy fails before initial fault can be repaired
- Time between fault and its repair is
 - ◆ *Window of Vulnerability (WOV)*

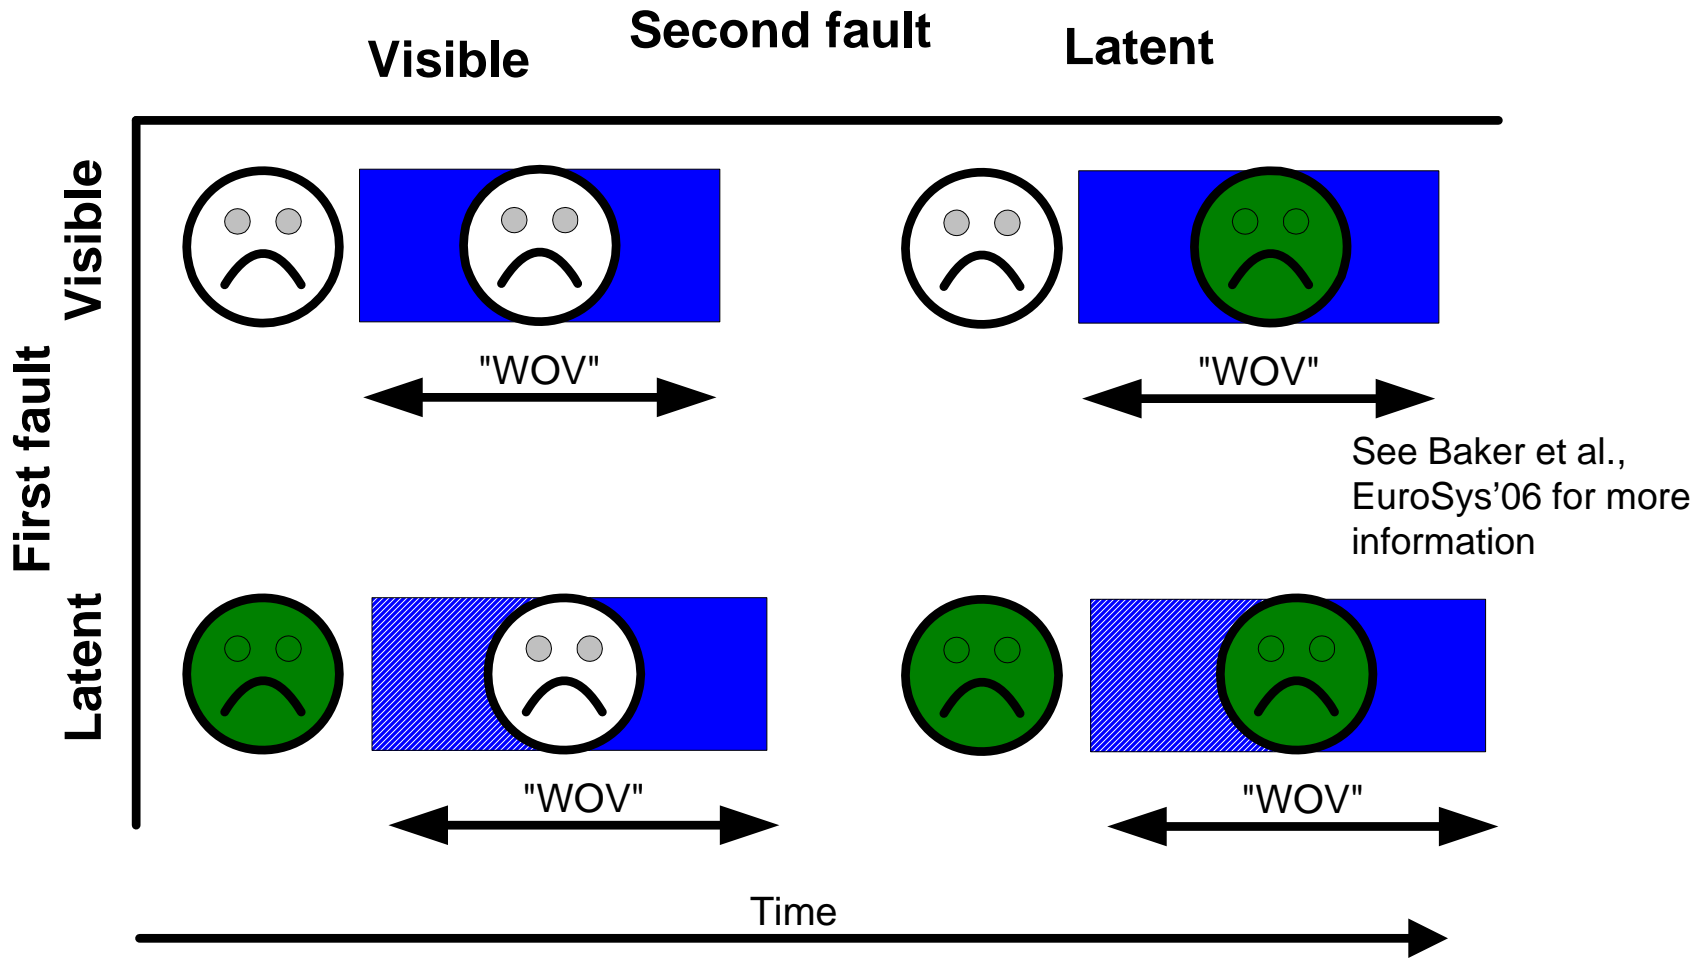
Window of vulnerability

Temporal overlap of faults



➤ Want detection time to be small


Data loss cases with 2 replicas



➤ Overall probability = sum of each case

Spatial overlap of faults

- Temporal overlap alone overstates likelihood of data loss



- Faults may be bits, sectors, files, disks, arrays, etc.
- If any two faults overlap, data is lost
- The smaller the faults, the less likelihood of overlap

Completing the model

- Multiply temporal and spatial probabilities
 - ◆ For each of the four loss cases
- Correlation: use multiplicative scaling factors for
 - ◆ Temporal correlation of faults
 - ◆ Spatial correlation of faults
- We also extend the model for further replication

Example using the model

- Shorter detection time helps how much?
- Portion of real archive (www.archive.org)
 - ◆ Monthly snapshots of web pages
 - ◆ 1.5 million immutable files
 - ◆ 1795 200GB ATA drives, “JBOD”
 - ◆ Mean time to visible (disk) failure: 20 hours
 - ◆ Almost 3 years of monthly file checksums
 - ◆ Mean time to latent fault 1531 hours
- See slide #15 for the results