Long Term Retention Work @ SNIA

Roger Cummings

Co-Chair, Long Term Retention TWG

Technical Director, CTO Office, Symantec

SNIA

About the SNIA LTR TWG

SNIA
Advancing storage & information technology

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.

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

Example: audited replicated archive

SNIA
Advancing storage & information technology

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

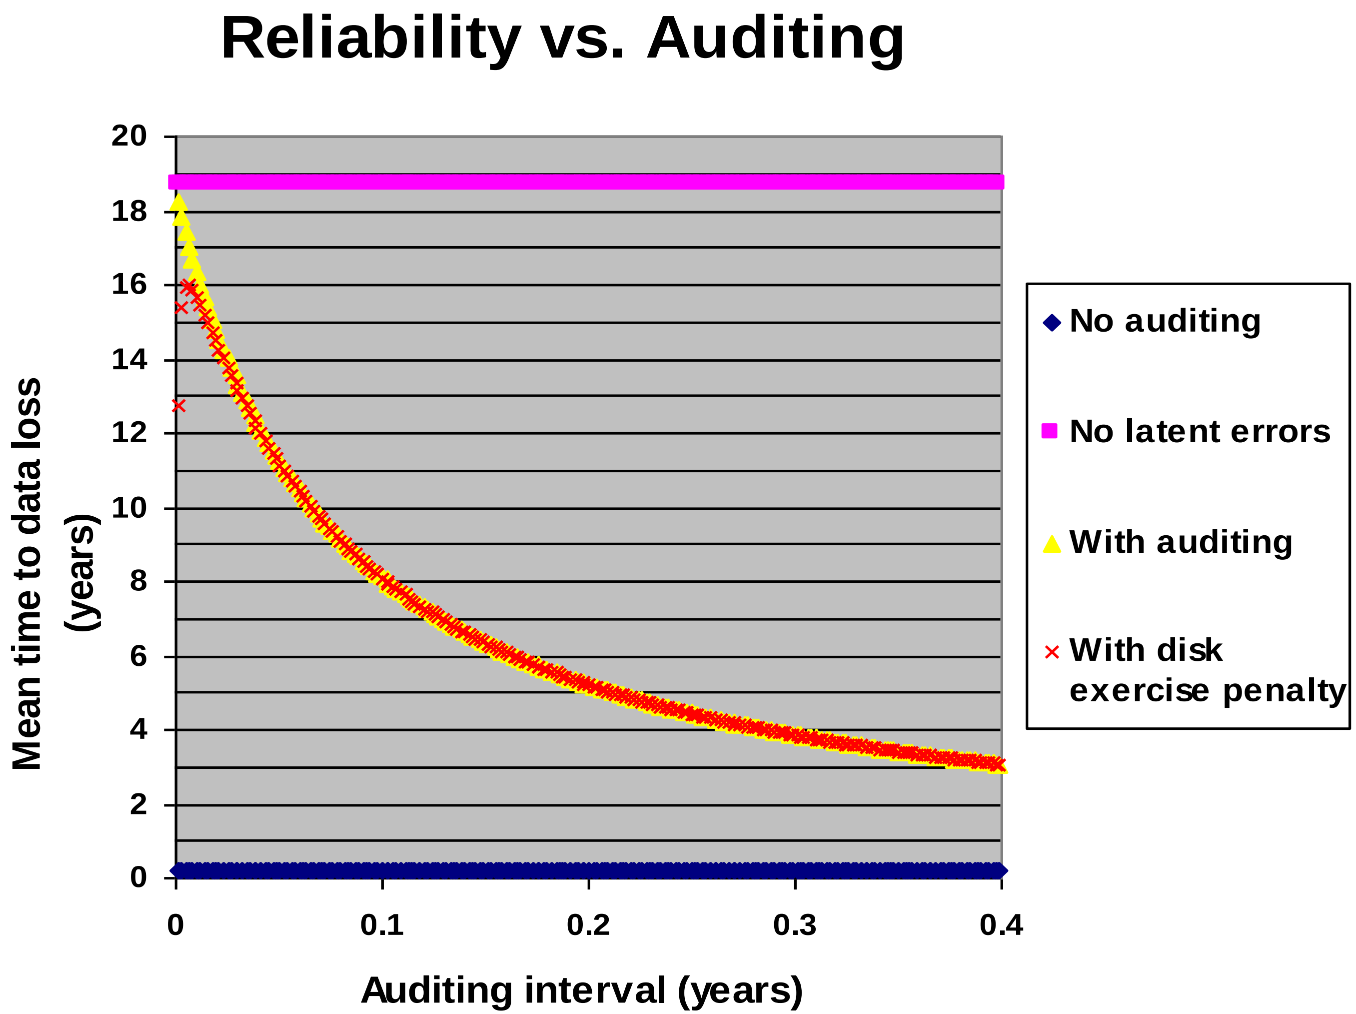

Best practices will vary over time

SNIA
Advancing storage & information technology

- 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

SNIA
Advancing storage & information technology

- 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

Preservation object storage formats

- We can't predict the future
  - Storage systems will change
  - Formats will change
  - Systems will fail
- Preservation objects need to be
  - **Self contained** – to ensure objects are complete
  - **Self describing** – so software can interpret it
  - **Extensible** – so it can meet future needs
- A preservation storage format must
  - Facilitate self contained, self-describing, and extensible object storage
  - Map to a wide variety of storage devices and technologies
  - Be resilient to failures and change

SIRF: logical container format

- SIRF – Self-Contained Information Retention Format
  - A format 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
  - Facilitates transparent migration for long-term preservation
    - Logical
    - Physical
- SIRF implementations will leverage other standards
  - Open Archival Information System (OAIS) ISO standard
  - Network Attached Storage (NFS/CIFS), XAM (Extensible Access Method), Object Storage (OSD)
  - Others

Problems SIRF addresses

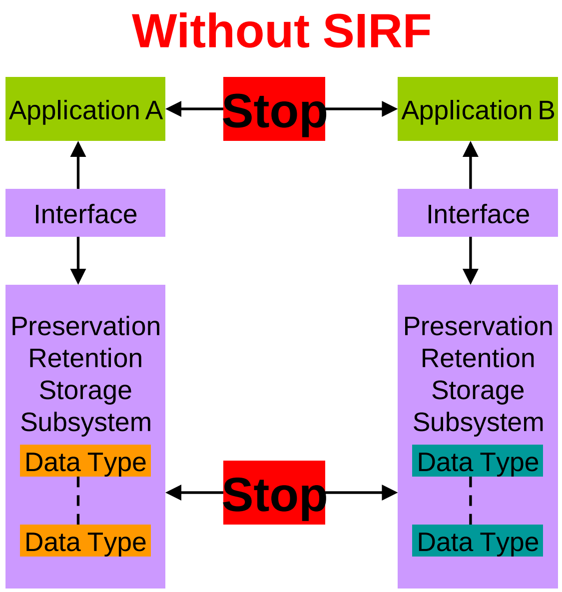

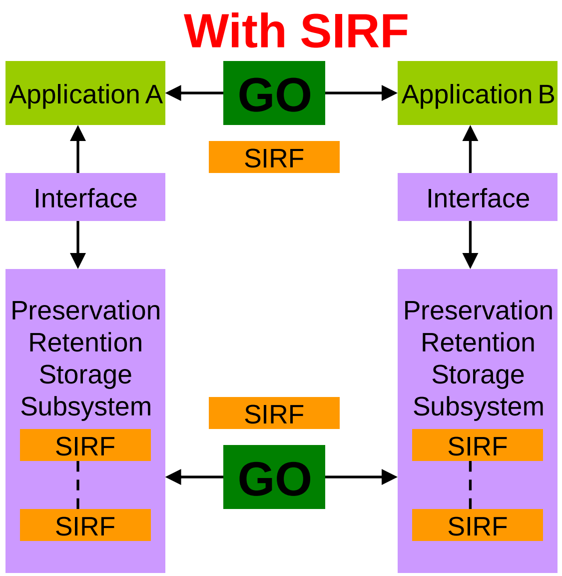

Cannot move cluster of preservation objects between systems without help	Can move cluster of preservation objects between systems by itself
Only original application that wrote the preservation objects can read and interpret them	Any SIRF compliant application can read and interpret preservation objects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

What is a Preservation Object?

- The raw data to be preserved plus additional embedded or linked metadata needed to enable the sustainability of the information encoded in the raw data for decades to come
  - The preservation object may be subject to physical and logical migrations
  - The preservation object may be dynamic and change over time
  - The updated preservation object is a new version of the original preservation object and must include audit log records of the changes that have occurred so authenticity may be verified
- An example of a preservation object is OAIS Archival Information Package (AIP)
  - An AIP includes recursive representation information that enables future interpretation of the raw data

Open Archival Information System (OAIS)

- ISO standard reference model (ISO:14721:2002)
- Provide fundamental ideas, concepts and a reference model for long-term archives
- Includes a functional model that describes all the entities and the interactions among them in a preservation system
- Archival Information Package (AIP) - a logical structure for the preservation object that needs to be stored to enable future interpretation

* OAIS Functional Model

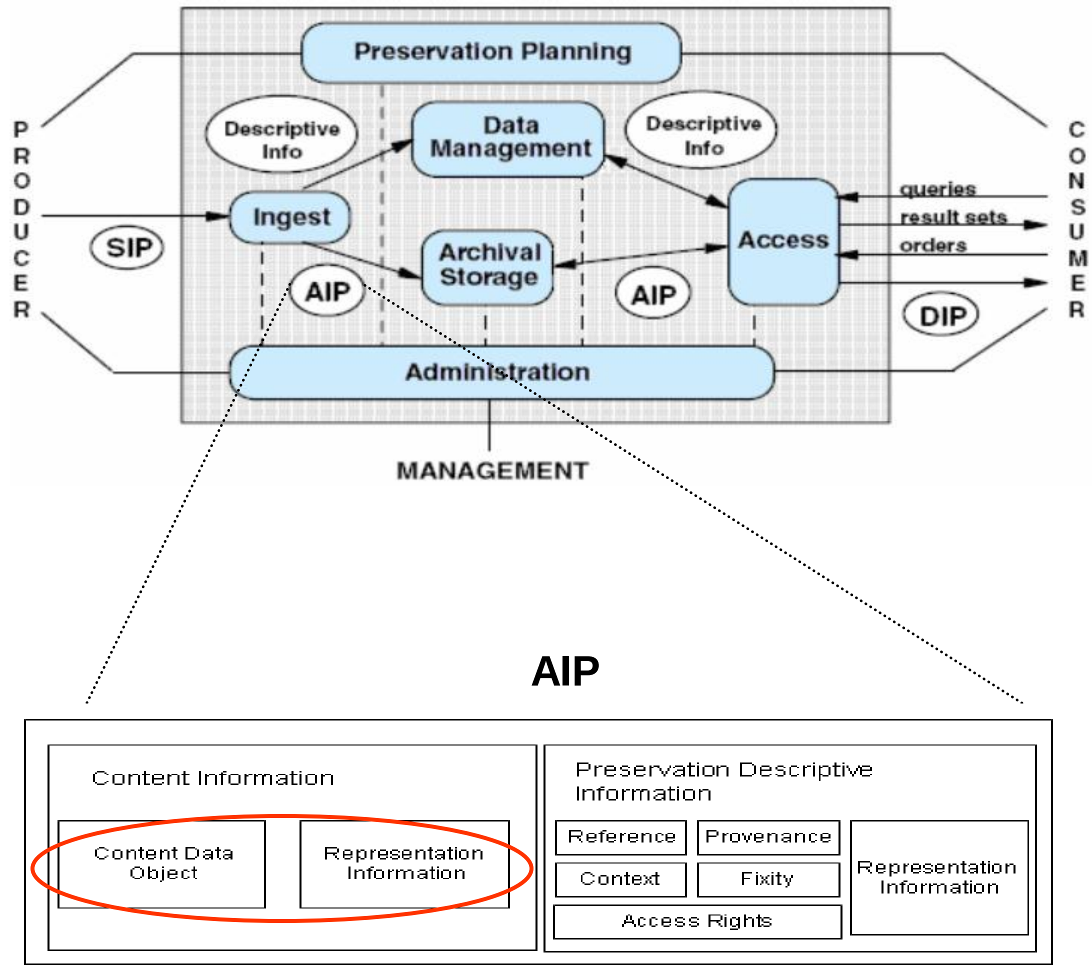

* Figure taken from the OAIS spec

SIRF Use Cases and Functional Requirements Methodology

- Define Actors involved in SIRF
- Define use cases and flows among the actors
  - Generic uses cases
    - Unlinked to specific type of data or application
    - Technological changes in the environment
  - Workload-based use cases
    - Specialized for concrete workloads
    - Additional non-technological changes in the environment
- For each use case, find the derived functional requirements
- Aggregate all functional requirements and map use cases to them
- Categorize the functional requirements
  - general requirements, format requirements, data model requirements, performance requirements, etc.
- Prioritize the functional requirements
  - Some of the requirements may conflict each other

Use Cases

- Generic Use Cases
  - Ingest and access with same application
  - Ingest and access with different applications
  - Ingest and access with different preservation services
  - Storage format is changed
- Workload-based Use Cases
  - eDiscovery
  - eMail archive
  - Consumer archive on cloud
  - BioMedical bank

SIRF Initial Requirements – General

- Media agnostic
  - Tape, disk, future media
  - Direct random access and serial access
  - Support mixture of storage technologies
- Vendor and Platform agnostic
- Support different standard storage technologies and interfaces e.g. NFS, CIFS, XAM
- Extensible
  - Support additional information which may be added in the future

SIRF Initial Requirements – Format

- Self-describing
  - The amount of a-priory information is small and can be acquired in stages
  - Interpretable by both humans and machines
  - Ability to do offline inspection
- Support self-contained data
  - Include means to represent internal links and cross references
  - Ability to reconstruct only from the container and any well defined external resources
- Support different SIRF formats preserved in an external registry
- Interoperability
  - Ability to migrate data between different systems without loss of information – data should be interpretable after migrations
  - Can be interpreted in the future
- Support methodology for verification of completeness and correctness

SIRF Initial Requirements - Preservation Object Data Model

SNIA
Advancing storage & information technology

- Support different data models for preservation objects
  - Support different object data models at one time
  - Support complex data structures like collections of objects
  - Support migrating objects from one data model to an alternative data model
- Can handle any proper data format for the raw data
  - No restrictions on file formats
- Enable keeping various versions of the same preservation object with their relations
  - References from new to existing preservation objects of the same version series
- There must be a persistent identifier for each preservation object
  - Include additional external identifiers

SIRF Initial Requirements - Performance

- Performance
  - Need to have good performance even for data that includes text and binaries
  - Support large objects e.g. web archiving objects, database archiving objects
  - Do not require complete scanning for access

- Enable parallel data migration
  - Enable parallel reads and writes

Current Status

SNIA
Advancing storage & information technology

- SIRF status
  - The SNIA LTR TWG is currently finishing up the requirements/use case definition phase for SIRF
    - More detailed use cases and storyboards
    - Formal technical requirements and documentation
  - Work on the specification itself is expected to begin late this year
- More information
  - SNIA Technical Working Groups (including the LTR TWG) – http://www.snia.org/tech_activities/workgroups
  - 100 year archive survey – http://www.snia.org/forums/dmf/knowledge/white_papers_and_reports/
  - SNW tutorial on Long Term Retention – http://www.snia.org/education/tutorials/2009/fall/
  - OAIS source doc – http://public.ccsds.org/publications/archive/650x0b1.pdf
  - XAM – http://www.snia.org/forums/xam/
  - NFS – http://www.ietf.org/dyn/wg/charter/nfsv4-charter.html

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!

Acknowledgements

- Simona Cohen, Sam Fineberg & Mary Baker of the SNIA LTR TWG generated much of this material
- The members of the LTR TWG have reviewed this material and provided constructive feedback
- Portions of this material have been used in SNIA tutorial @ StorageNetworkingWorld conferences worldwide

Advancing storage & information technology

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

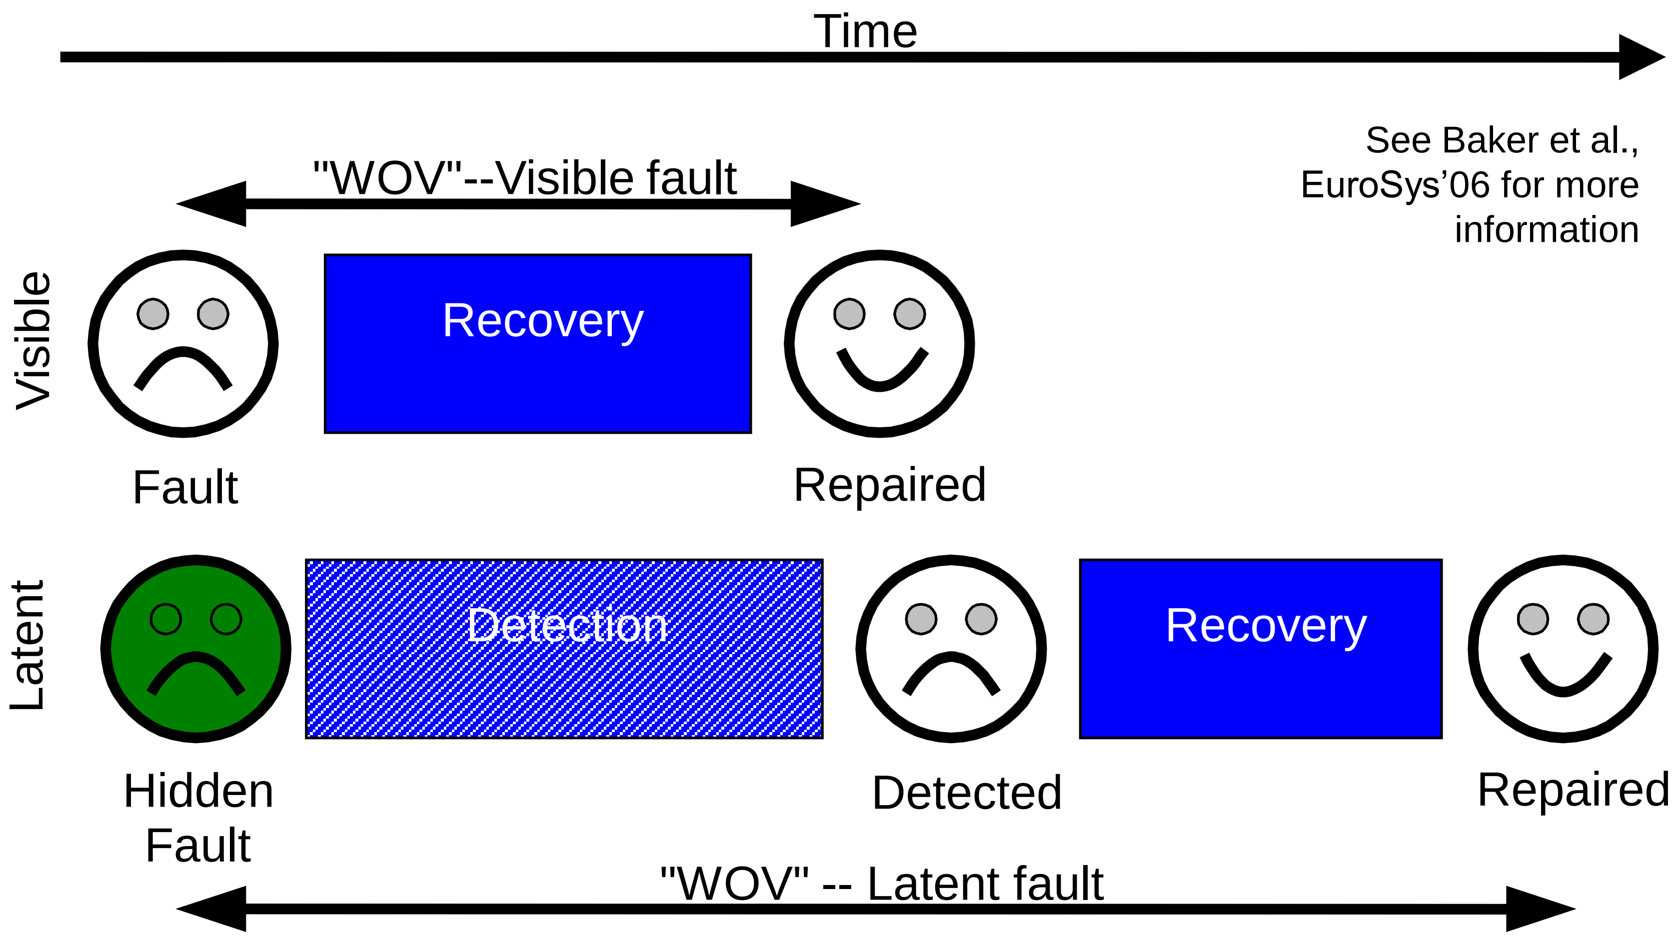

- Want detection time to be small

Data loss cases with 2 replicas

SNIA
Advancing storage & information technology

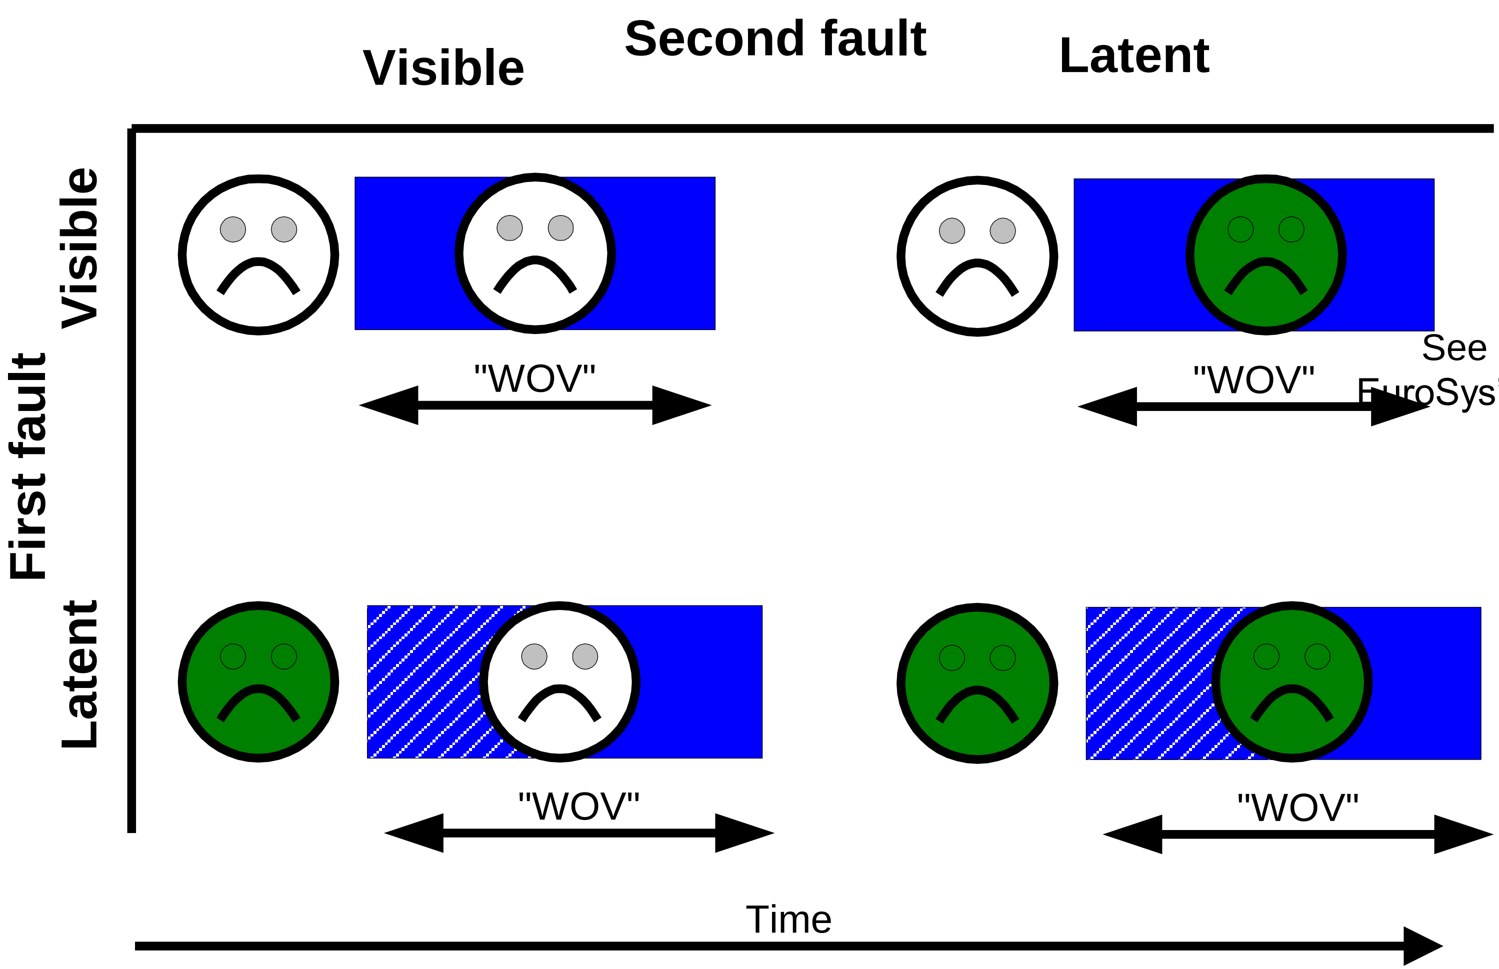

See Baker et al., EuroSys'06 for more information

- Overall probability = sum of each case

Spatial overlap of faults

- Temporal overlap alone overstates likelihood of data loss

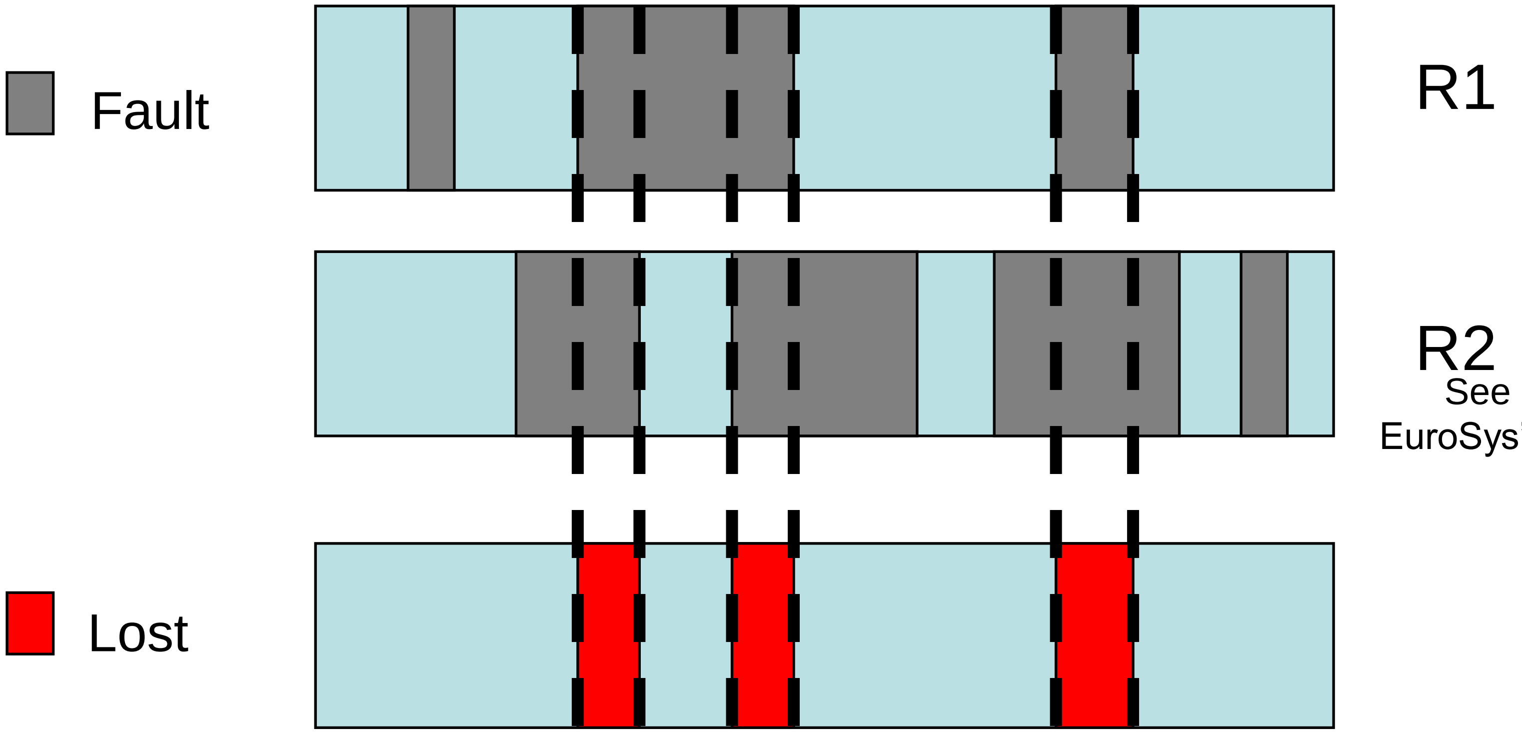

See Baker et al., EuroSys'06 for more information

- 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

Advancing storage & information technology

- 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 #4 for the results