HP WORLD 2004
Solutions and Technology Conference & Expo

Defending data against disasters

Dr. Kimberly Keeton

Senior Research Scientist
Hewlett-Packard Labs

Session outline

- Motivation
- Overview of data dependability designer
- Your feedback (interactive)
- Conclusions

Motivation

- Since "disasters" happen, it's only wise to protect against them
- High cost of unavailability ($/hour downtime):
  - Brokerage operations $6.4M
  - Credit card authorization $2.6M
  - Ebay (1x22-hr outage) $225K
  - Amazon.com $180K
  - Airline reservation center $89K

Source: InternetWeek 4/3/2000 + *Fibre Channel: A Comprehensive Introduction*, R. Kembel 2000, p.8. "...based on a survey done by Contingency Planning Research."

- High cost of data loss: Gallup poll: 100MB == $1M

Source: "The Data Recovery Solution," white paper by OnTrack Data Recovery, Inc., 1998, available from http://www.ontrack.com.

Motivation

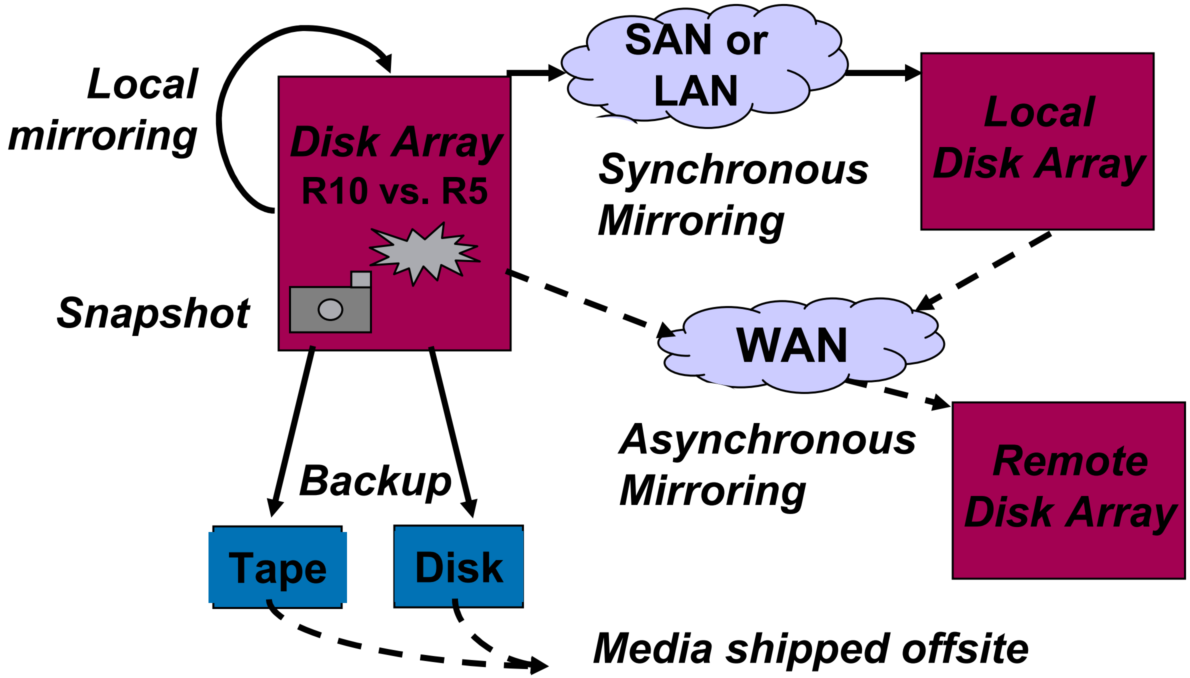

- Determining how to meet dependability goals is hard
  - Increasing number of data protection mechanisms
  - Lots of configuration parameters
- Today's design techniques: manual, ad hoc approaches
  - Insufficient tools support for examining wide range of candidate designs
  - Current designs are likely conservative
  - Only qualitative understanding of design dependability

Where do customers want help?

- Scenario 1: customer wants to understand whether their configuration meets their needs
- Scenario 2: customer's IT-savvy sysadmin needs help justifying her technology choices to business management
- Scenario 3: customer hires HP for business impact assessment; what's the best design for their needs?
- Scenario 4: customer needs help understanding how business requirements and design choices influence solution cost

Our research: data dependability designer

- Solver to automatically design basic data dependability solutions
- Evaluate business impact of a particular solution
  - *Outlay* costs for equipment, facilities, service
  - *Penalty* costs for recovery time and recent data loss
- Pick best solution for specified inputs
  - Business needs
  - Workload requirements
  - Failure scenario
- Explore sensitivity of solution choice and cost to input specification

"Designing for disasters", K. Keeton, C. Santos, D. Beyer, J. Chase, and J. Wilkes. *Proc. 3rd USENIX Conference on File and Storage Technologies (FAST)*, March 2004.

"A framework for evaluating storage system dependability," K. Keeton and A. Merchant. *Proc. Intl. Conference on Dependable Systems and Networks (DSN)*, June 2004.

Designer at a glance

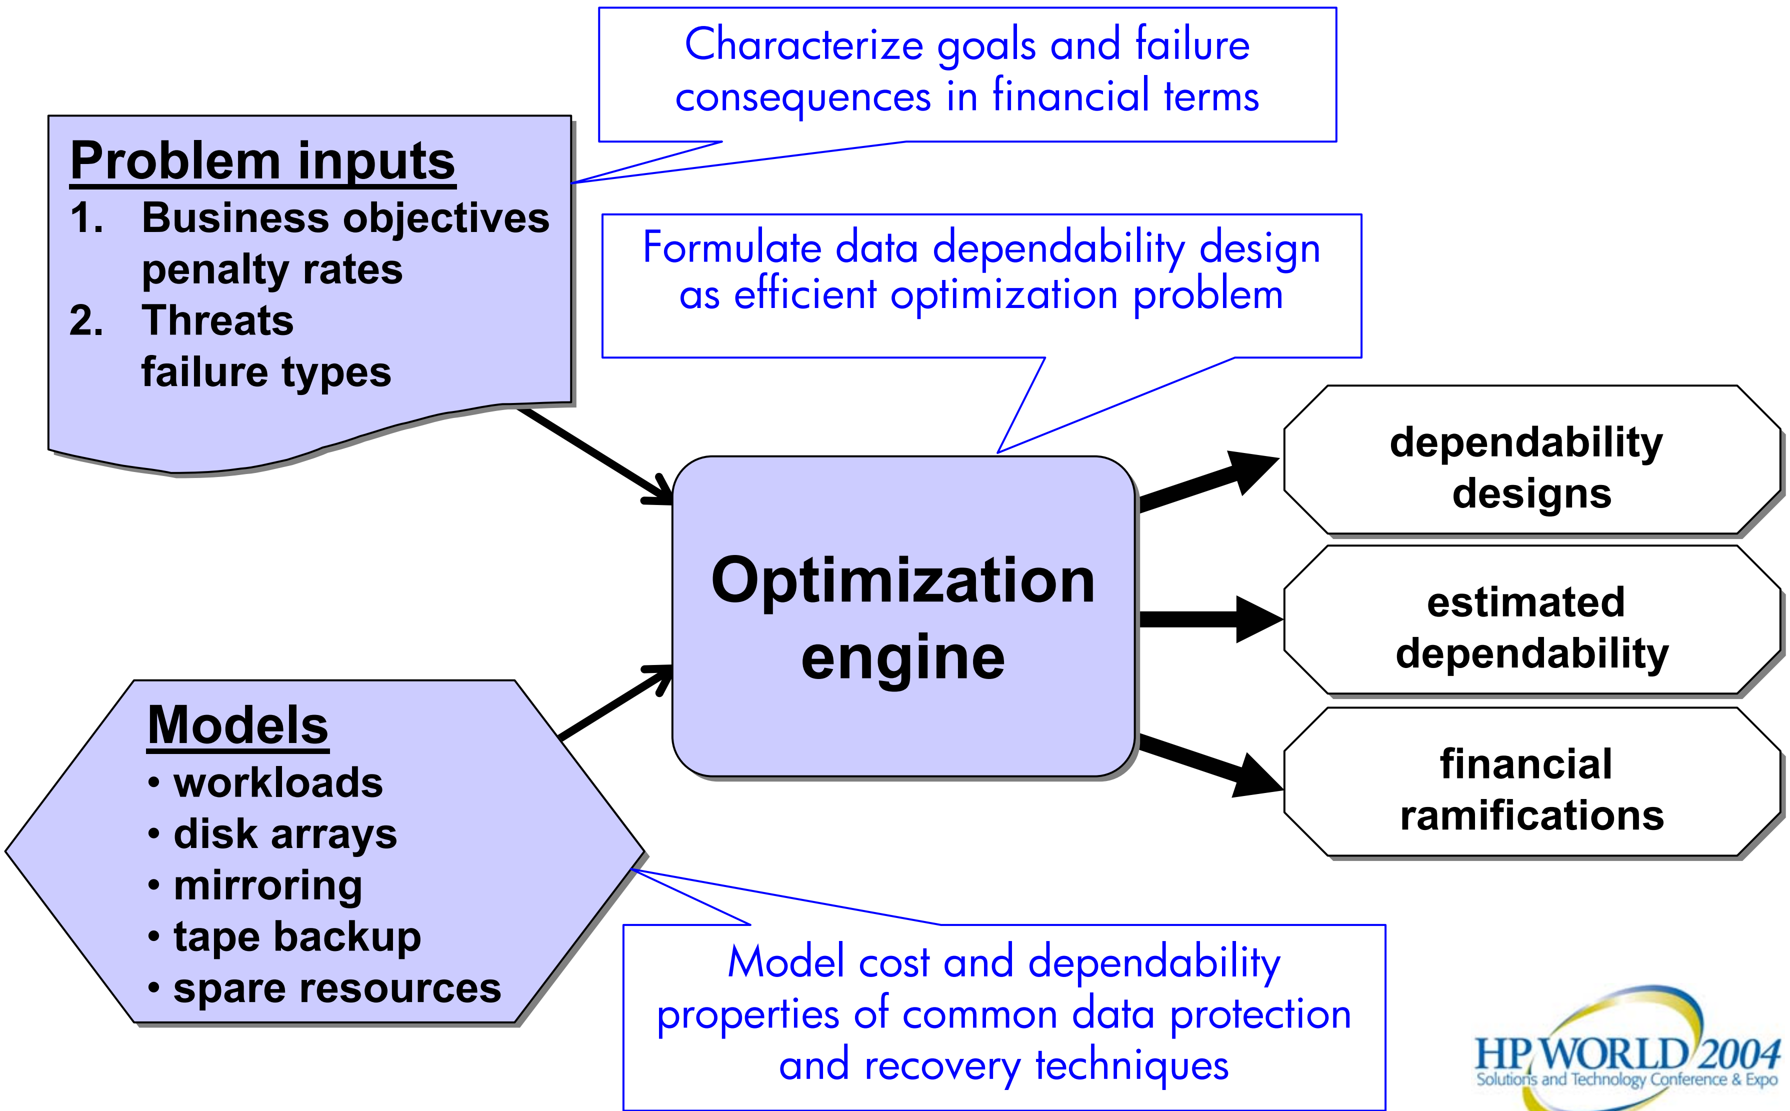

HP WORLD 2004
Solutions and Technology Conference & Expo

Benefits for HP's customers

- Ability to assess dependability of customer configurations
- Solutions that are potentially better matched to requirements
- Significantly reduced time to identify appropriate solutions
- Better customer understanding of potential solutions and their behaviors
- Better customer understanding of financial impacts of solution dependability

Goals for this session

- Provide overview of automated data dependability designer
  - What does the dependability designer do?
  - How does it work?
  - What questions can it help answer?

- Gather your feedback (interactive)
  - How do you design dependable storage systems today?
  - How much can you tell us about your requirements?
  - What do you need to inform your decision-making?
  - How would you want to use a tool like this?

Dependability designer overview

Data protection techniques

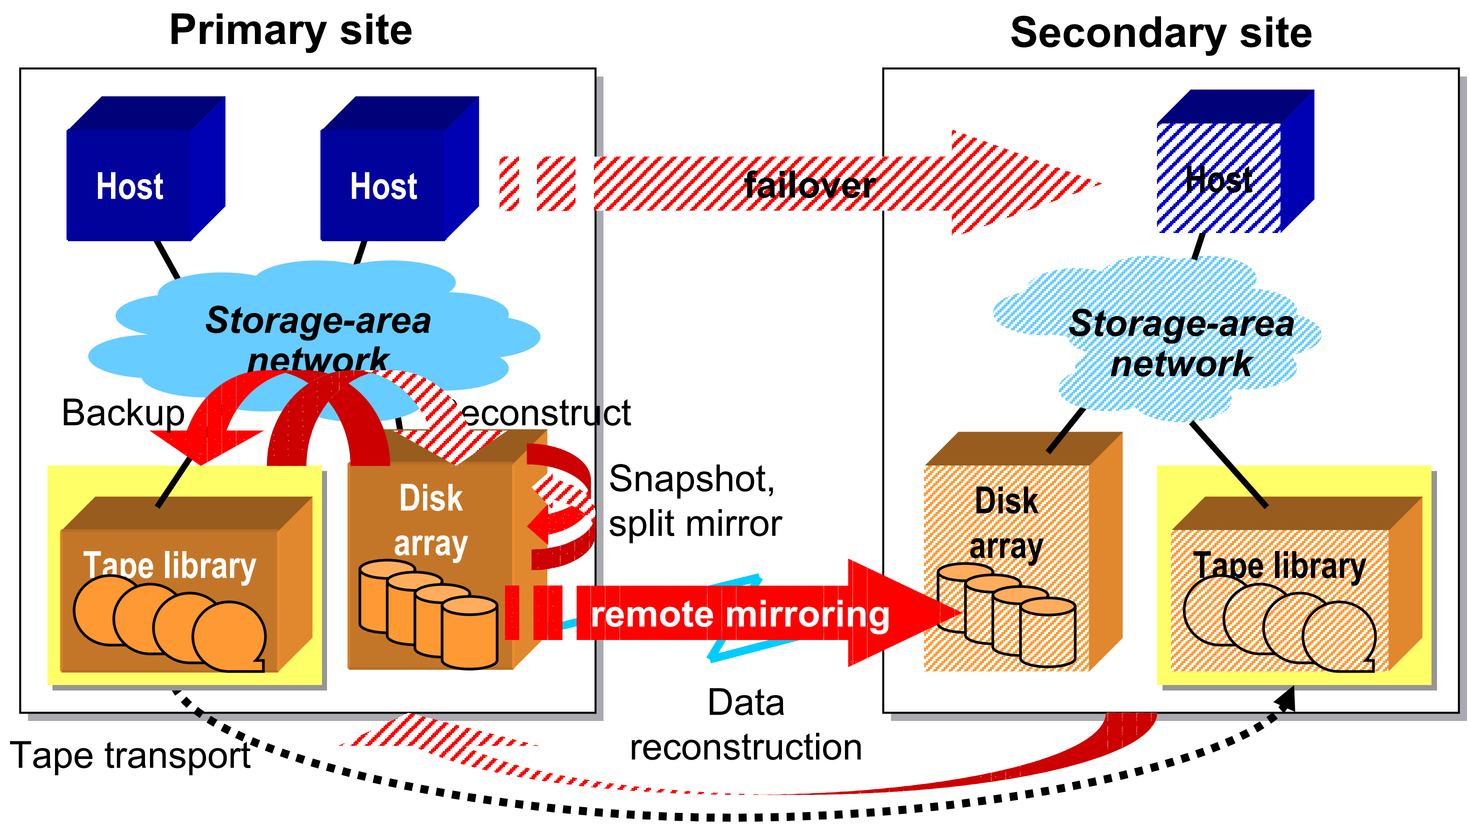

- Primary copy protected by one or more secondary copies
  - Local, regional, remote
- Secondary copy techniques modeled
  - Intra-array mirroring: snapshots, clones/split mirrors
- Secondary copy techniques (cont.)
  - Remote mirroring: sync, async, async with batching
  - Tape backup and vaulting
  - Failover vs. reconstruction
  - Resource sparing: hot vs. unconfigured, dedicated vs. shared

Tape backup and vaulting

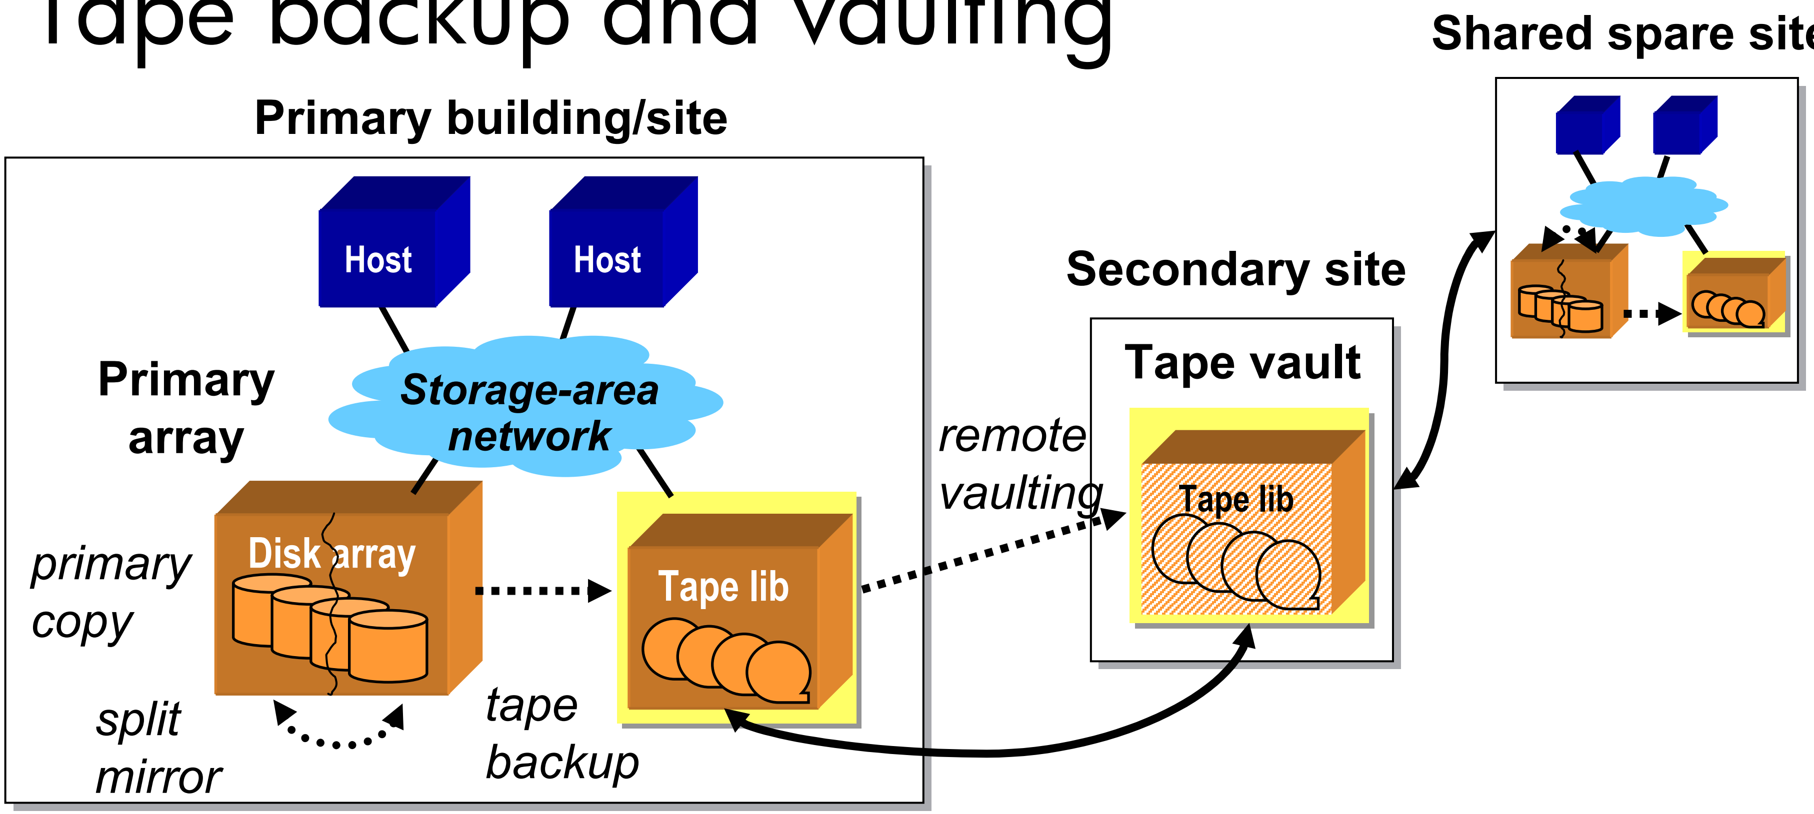

- Backup configuration questions:
  - How long between successive backups?
  - How often to do full vs. incremental backups?
  - How long should backup window be?
  - How long to keep backups?
- Vaulting configuration questions:
  - How often to ship tapes offsite?
  - How long to delay before shipping?
  - What to ship offsite?

Remote mirroring

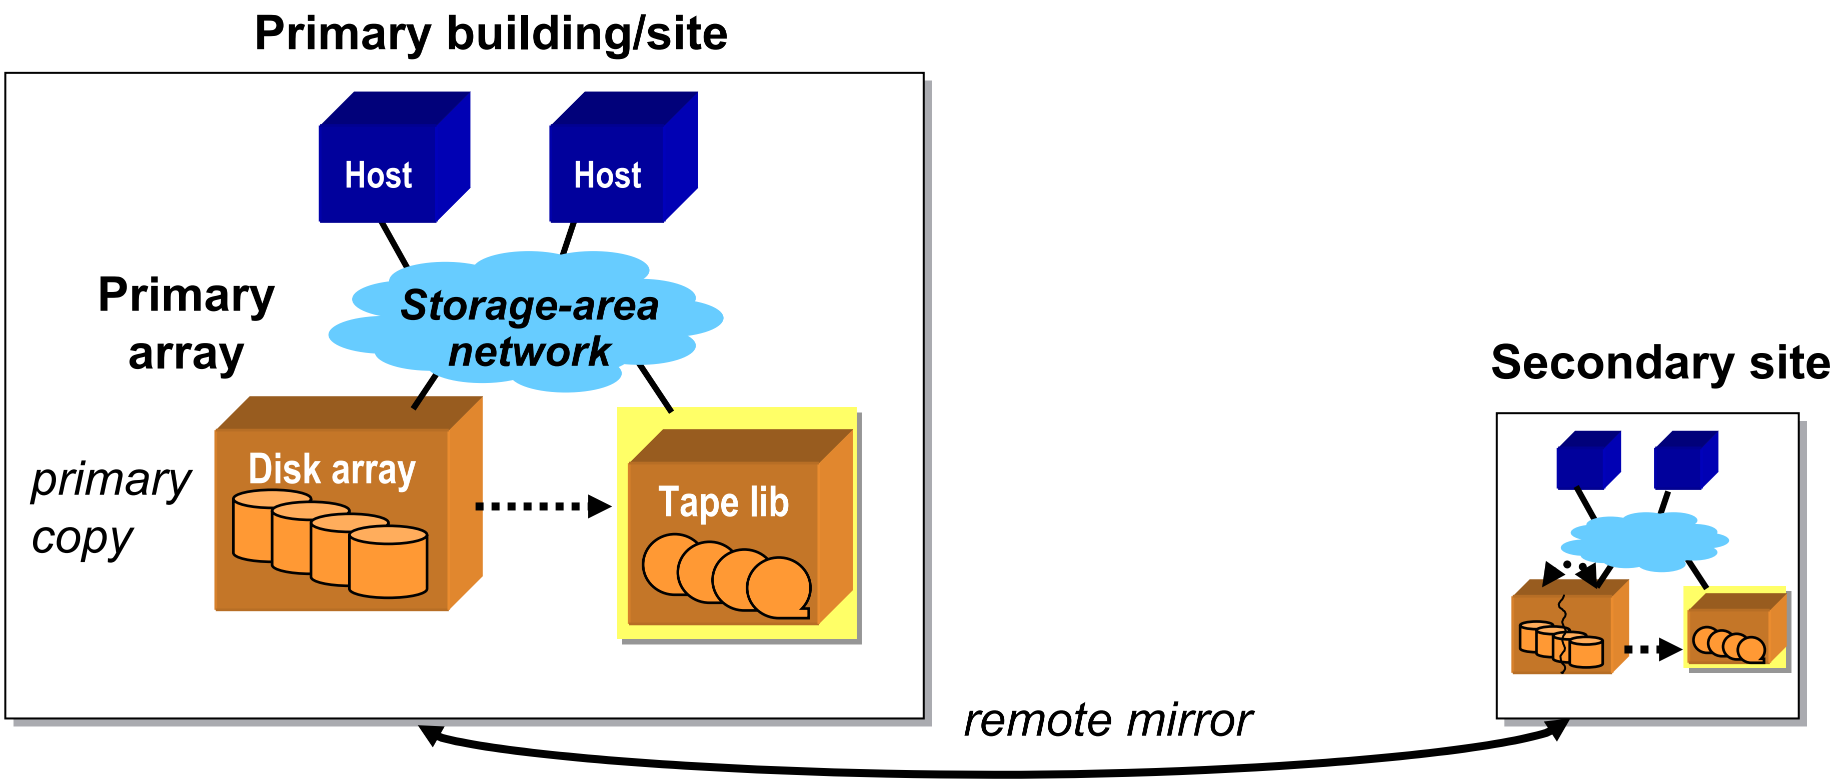

- Remote mirroring configuration questions:
  - What protocol to use – synchronous or asynchronous?
  - If asynchronous batch protocol, how long to coalesce updates?
  - How many network links to use?

Determining the right solution

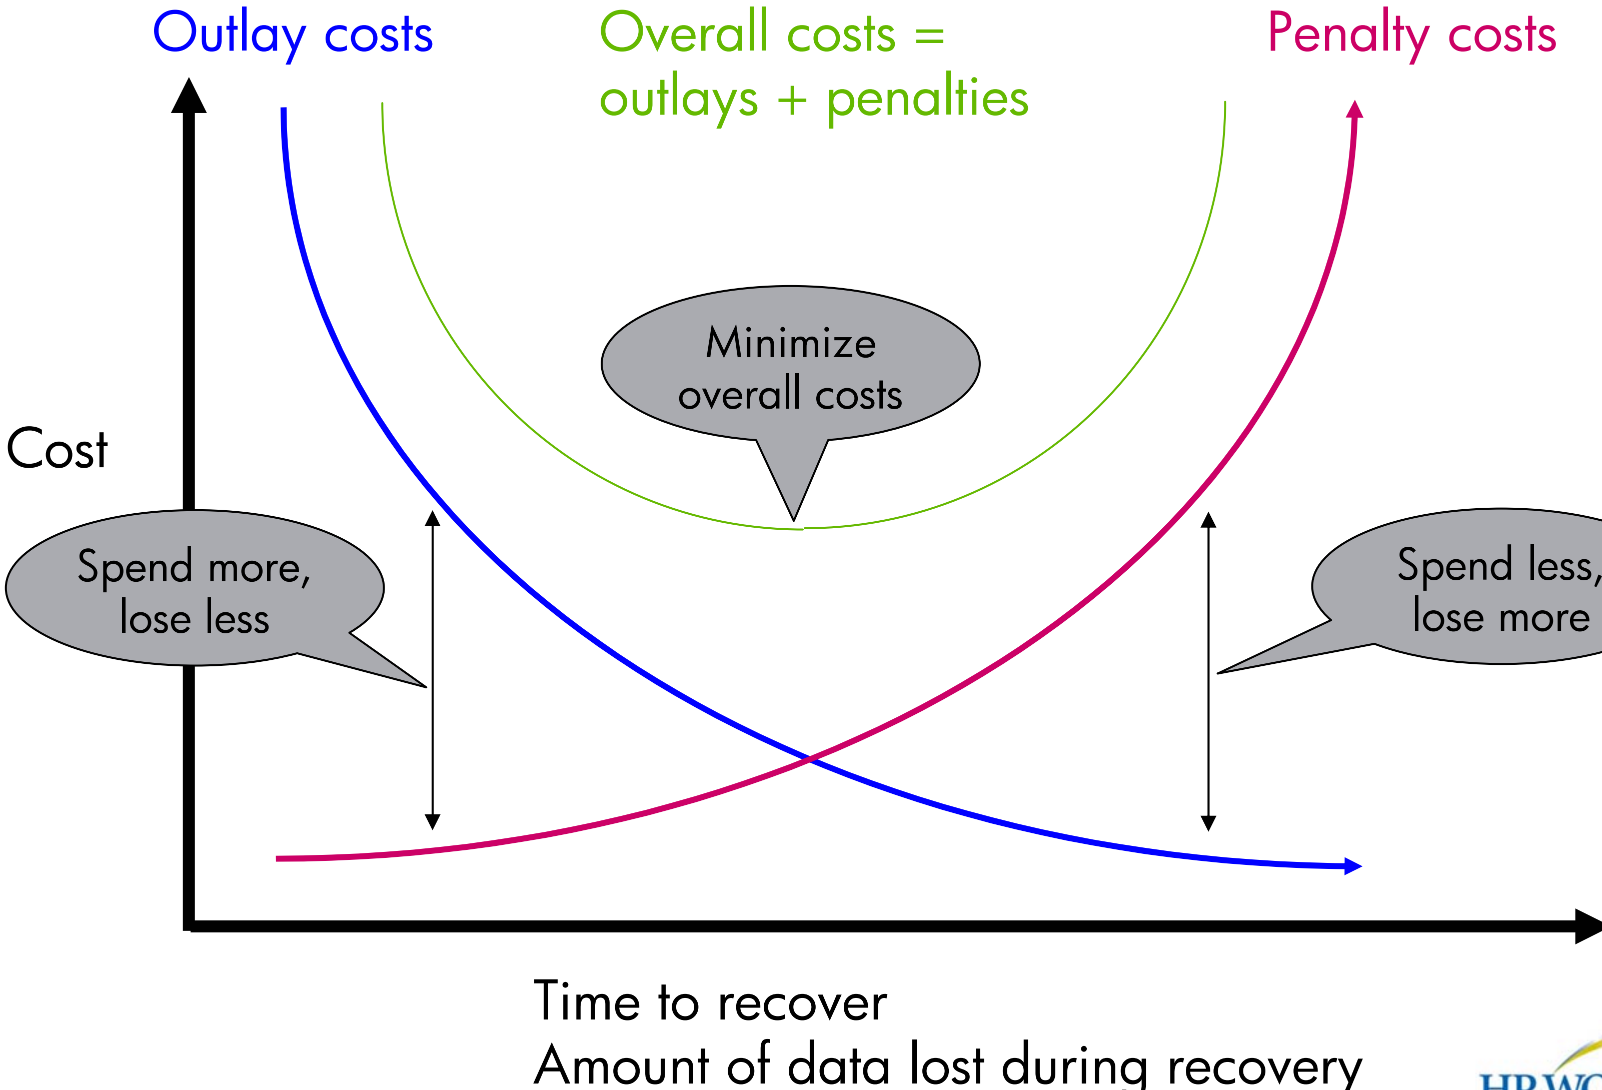

Dependability as optimization problem

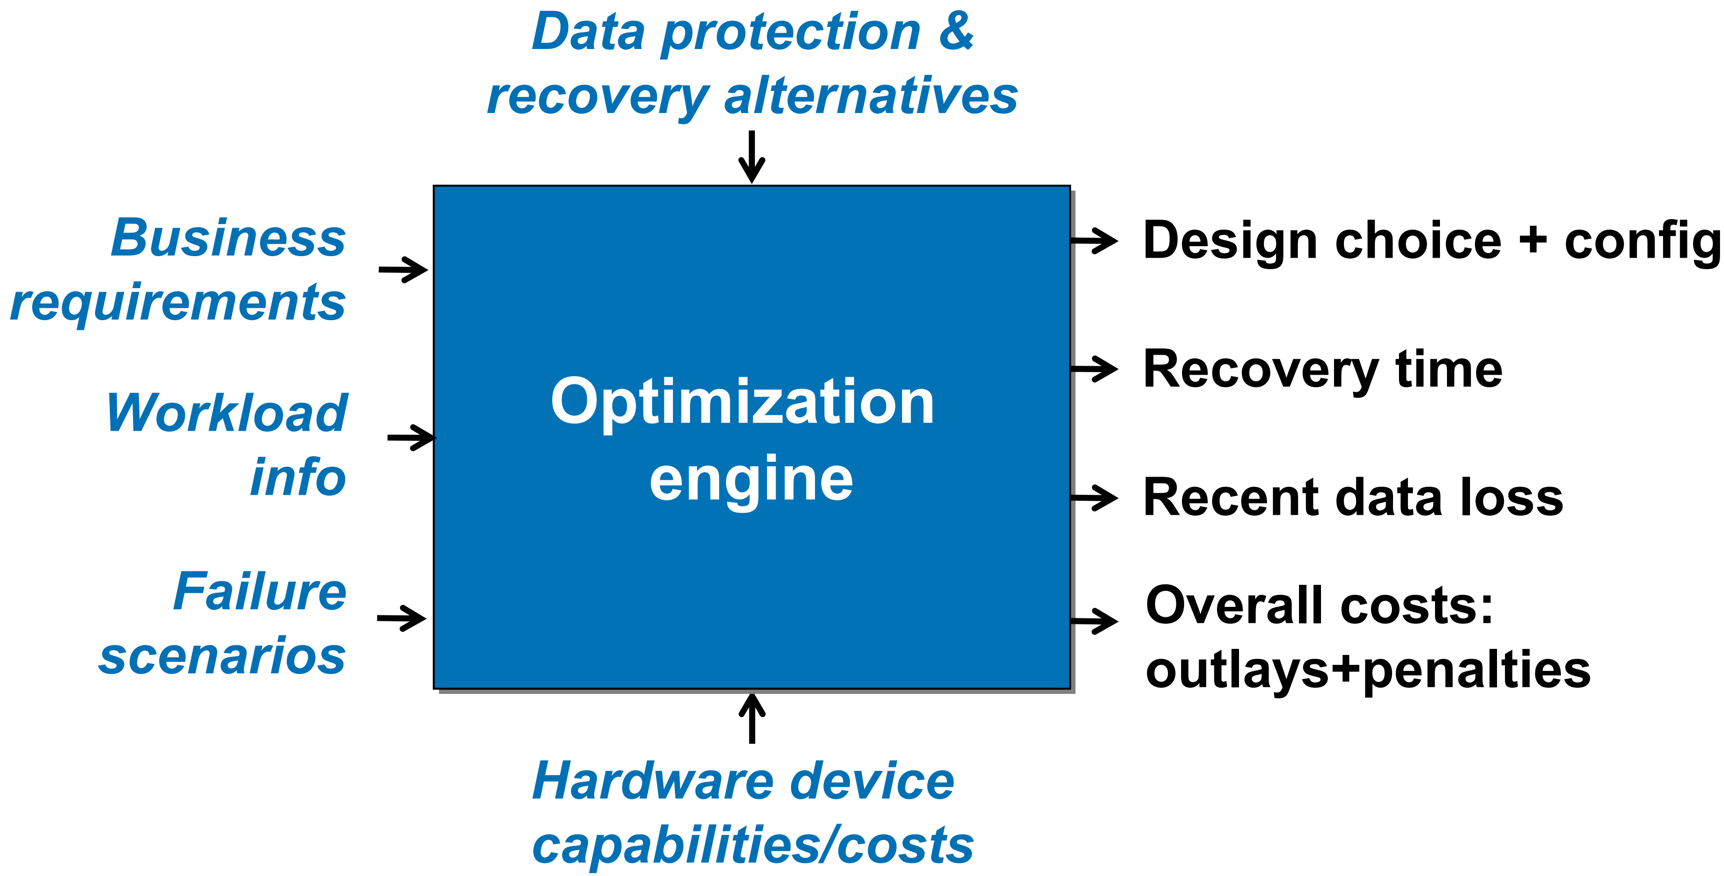

- Objective function
  - Minimize overall business cost = outlays + penalties

Business requirements: penalty rates

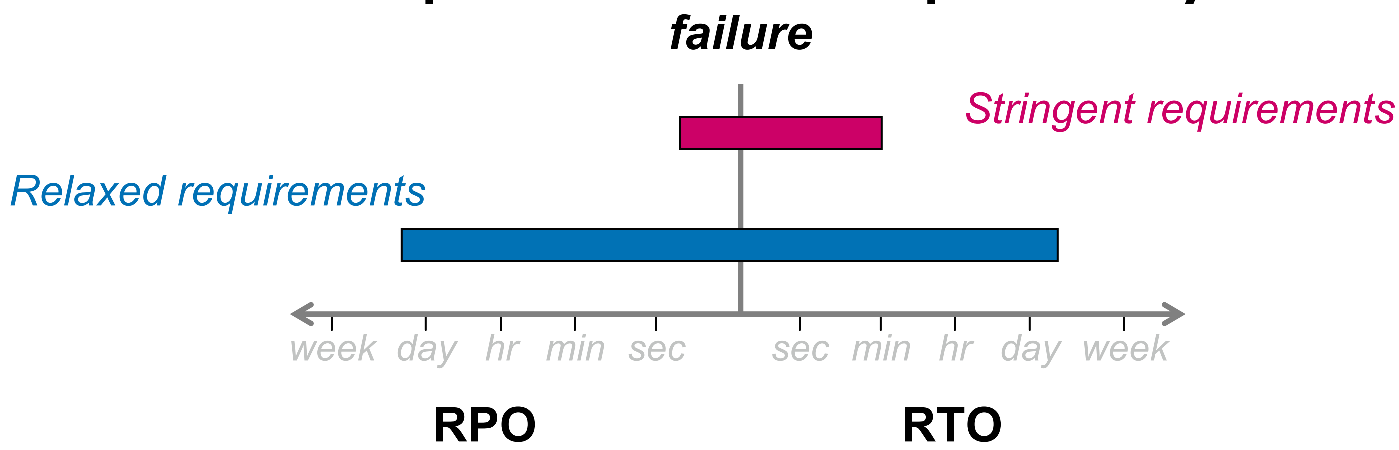

- Recovery time objective (RTO):
  - How long before the system is back up?
- Recovery point objective (RPO):
  - How much recent data can the system discard?

Business requirements: penalty rates

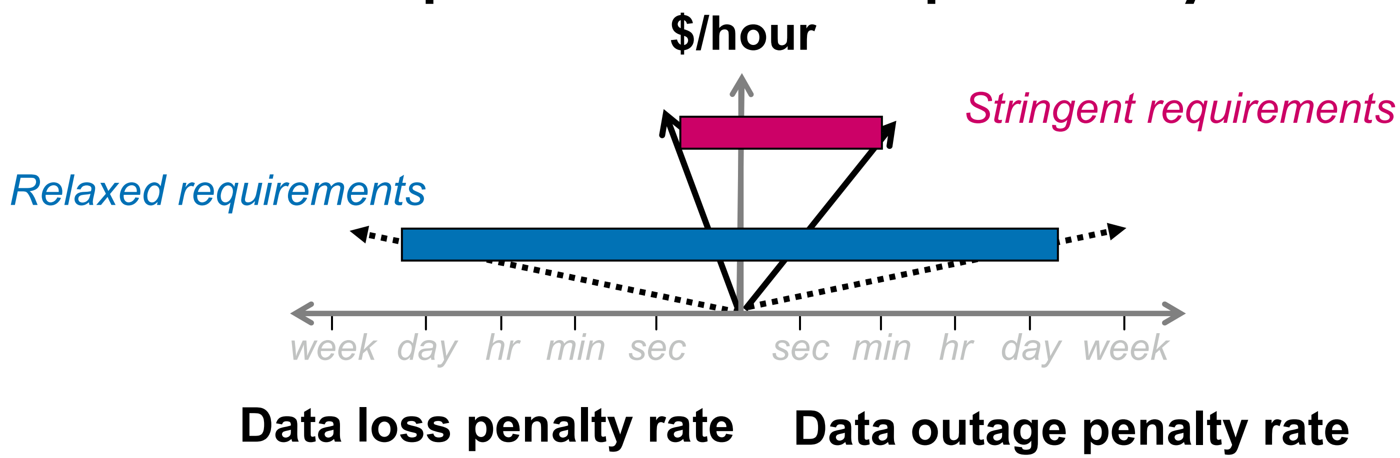

- Recovery time objective (RTO):
  - How long before the system is back up?
- Recovery point objective (RPO):
  - How much recent data can the system discard?
- Penalty rate model
  - Data loss penalty rate ($/hour)
  - Data outage penalty rate ($/hour)

Workload requirements

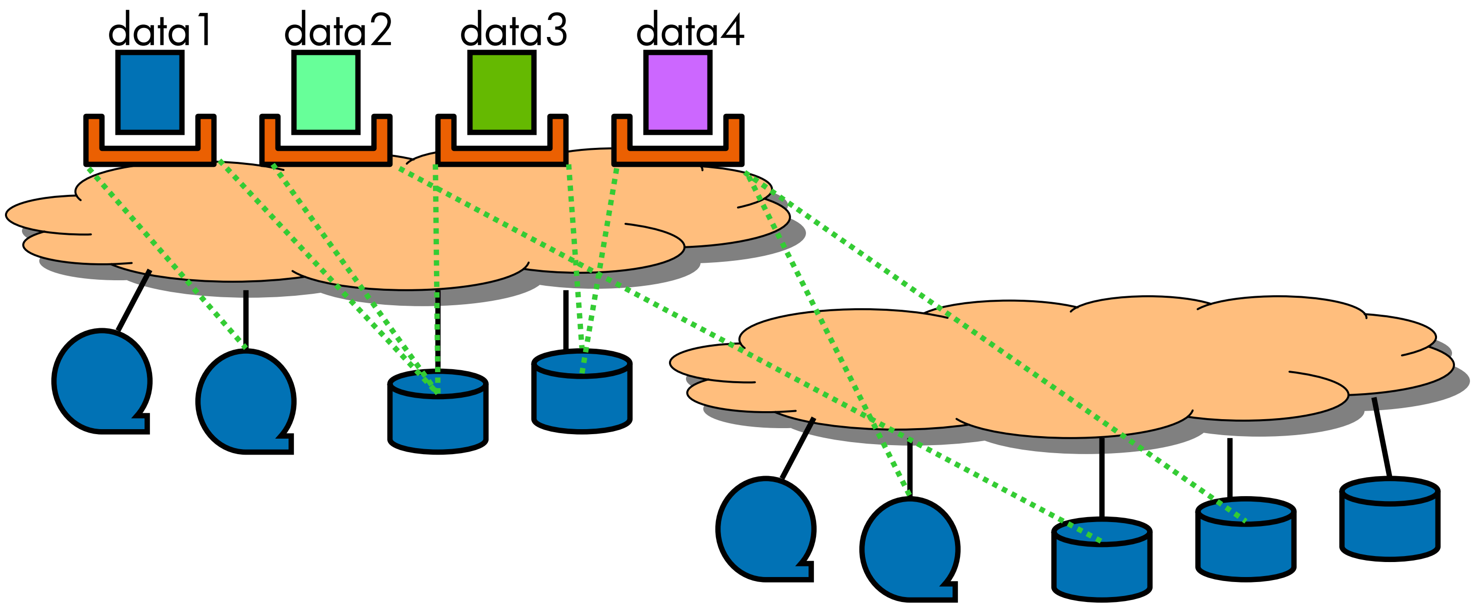

- Useful workload characteristics (per data object)
  - Capacity
  - Access rates
  - Update rates (both with and without overwrites)
  - Burstiness

Failures

- Our focus: recovery from primary copy loss due to:
  - "Container" failure (ex: primary array, primary site)
  - User or software error
- Recently written data may be more vulnerable
- Compute expected penalties based on specified failures and their relative frequencies of occurrence

Designer case studies (FAST '04)

- Evaluation of existing designs
- What if scenario analysis
- Automated design choices
- Dependability choice exploration
  - System dependability
    - Recovery time
    - Recent data loss
  - Overall costs

Evaluation of existing designs

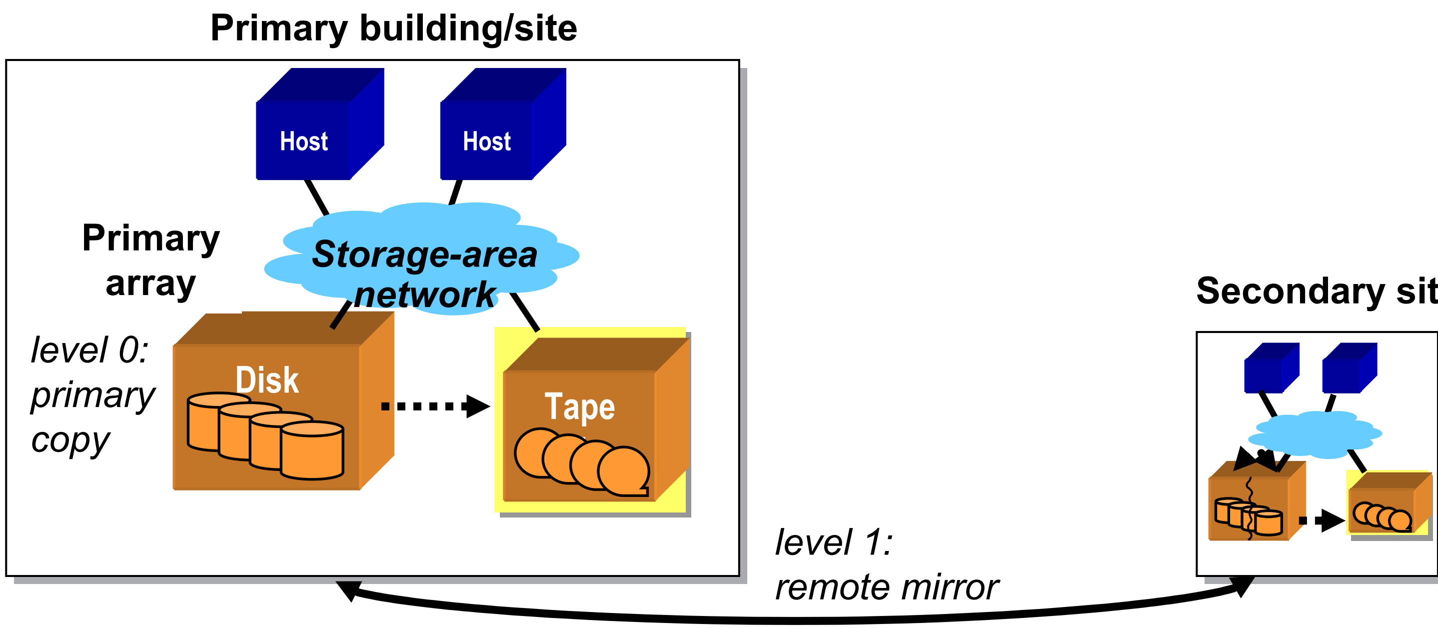

- Design: asynchronous mirroring, single T3 link
- Business requirements:
  - $20K / hour downtime
  - $20K / hour recent data loss
- Failure scenario:
  - One site disaster per year
- Workgroup file server workload:
  - Capacity: 1.36 TB
  - Average (non-unique) update rate: 799 KiB/s
  - Peak:average bandwidth burst multiplier: 10X
  - Batched unique update rate:
    - <1 min, 727 KiB/s> …
    - <24 hr, 317 KiB/s>

Evaluation of existing designs

- System dependability
  - Recovery time: 72 hours
  - Recent data loss: 2 minutes

- Financial ramifications
  - Outlay costs (annualized): $501K
  - Penalty costs:
    - Data outage penalties: $1.44M
    - Recent data loss penalties: $730
  - Overall costs: $1.95M

Asynchronous mirroring "what if"

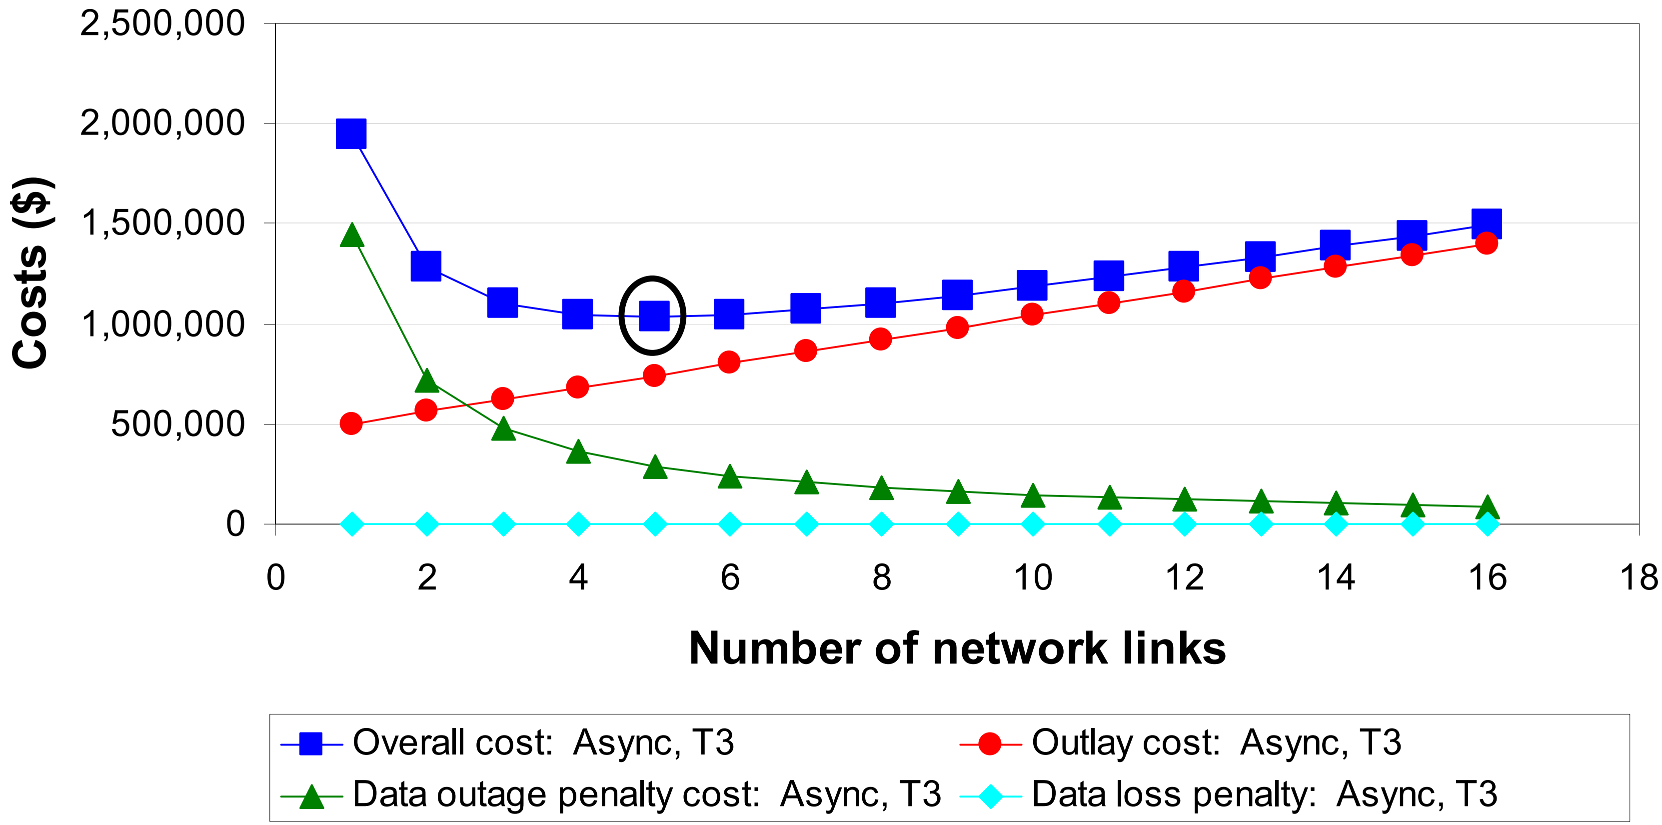

- Asynchronous mirroring with T3 links
- Minimal overall cost "sweet spot" at five links
  - Fewer links: outage penalties dominate
  - More links: outlay costs dominate

Automated design choices

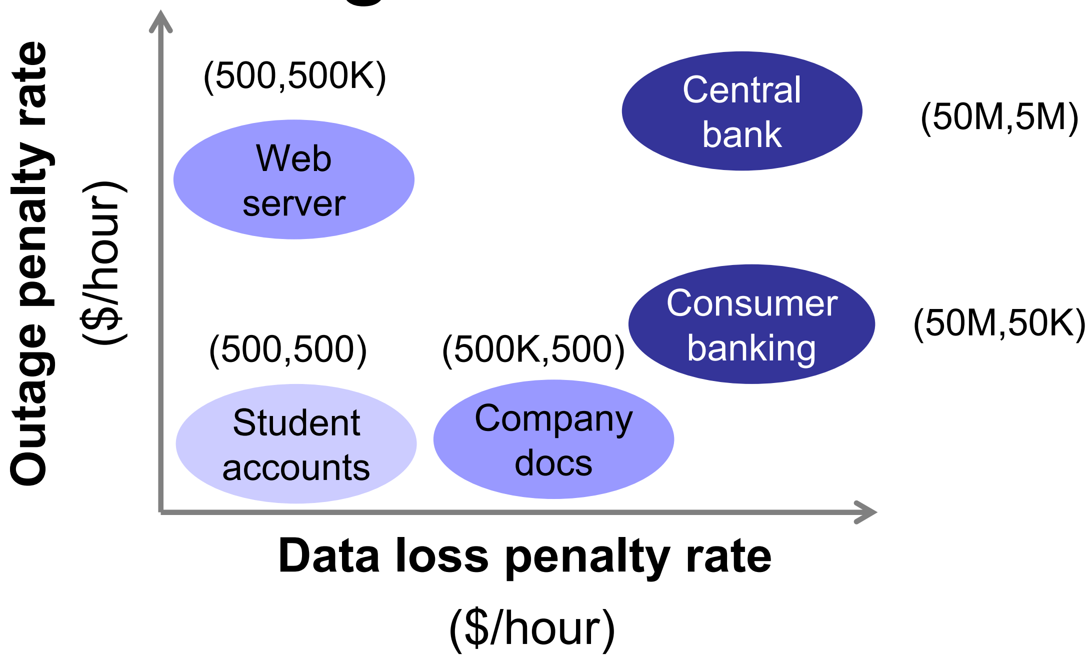

- Experimental design
  - Penalty rates for five different industry segments
  - Same workload ([cello2002] workgroup file server)
  - Annualized outlay costs, one site disaster per year
  - Solver determines best design

Automated design choices

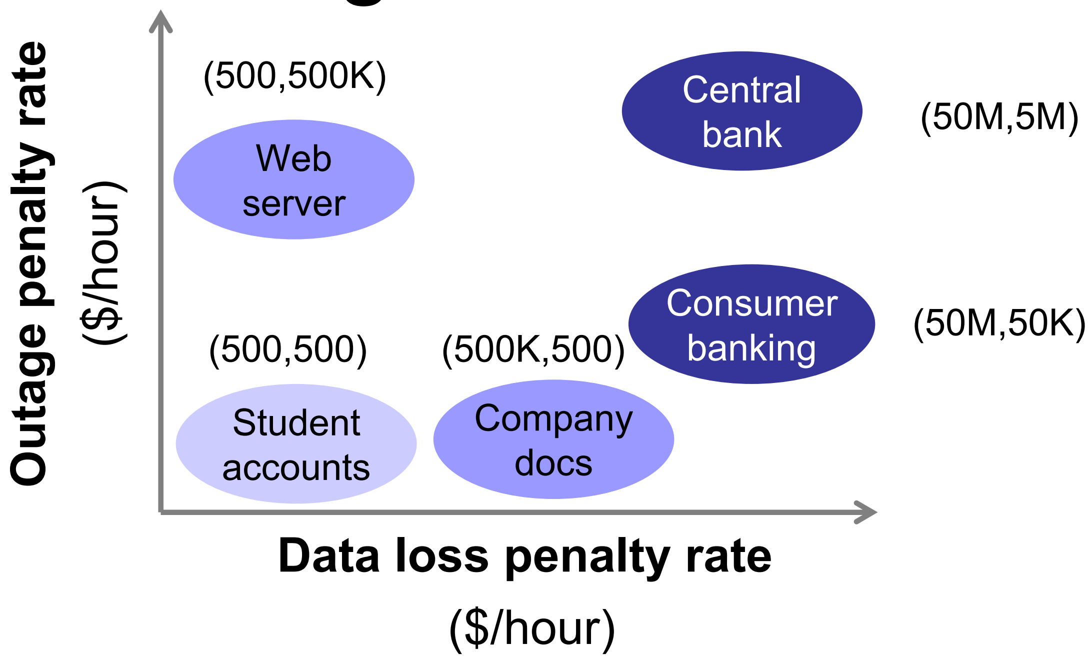

Industry segment	Automated design choice
Student accts	Backup + 12-hr win + 1 drive + no spares
Company docs	Async + 1 link + recovery + no spares
Web server	Batched async + 1 link + failover
Consumer banking	Sync + 2 links + failover
Central bank	Sync + 2 links + failover

Design space exploration

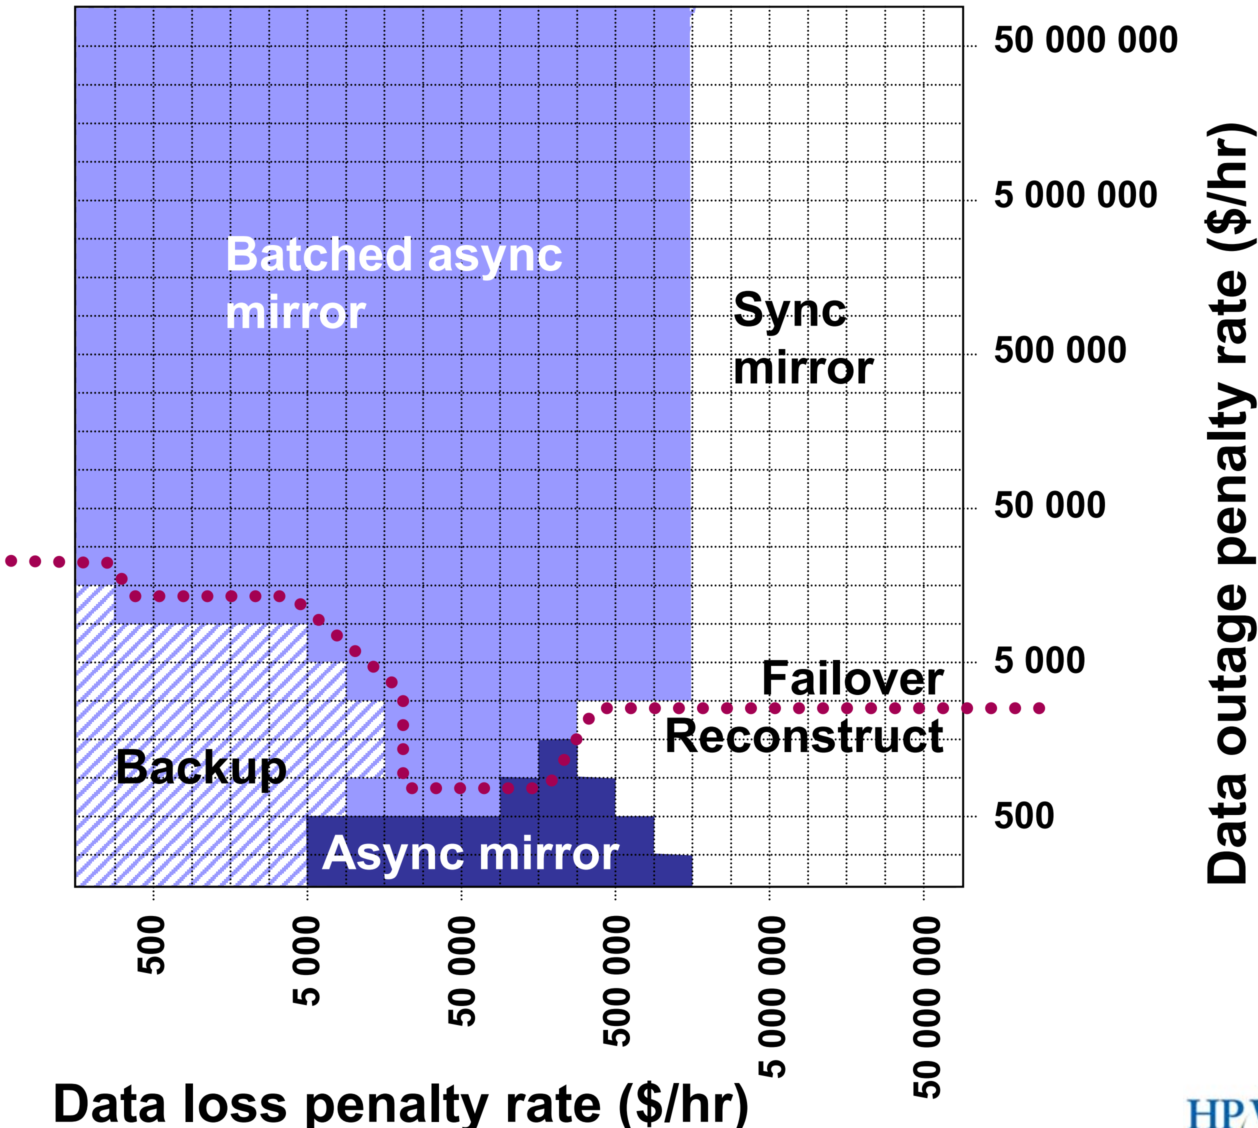

Design recovery time

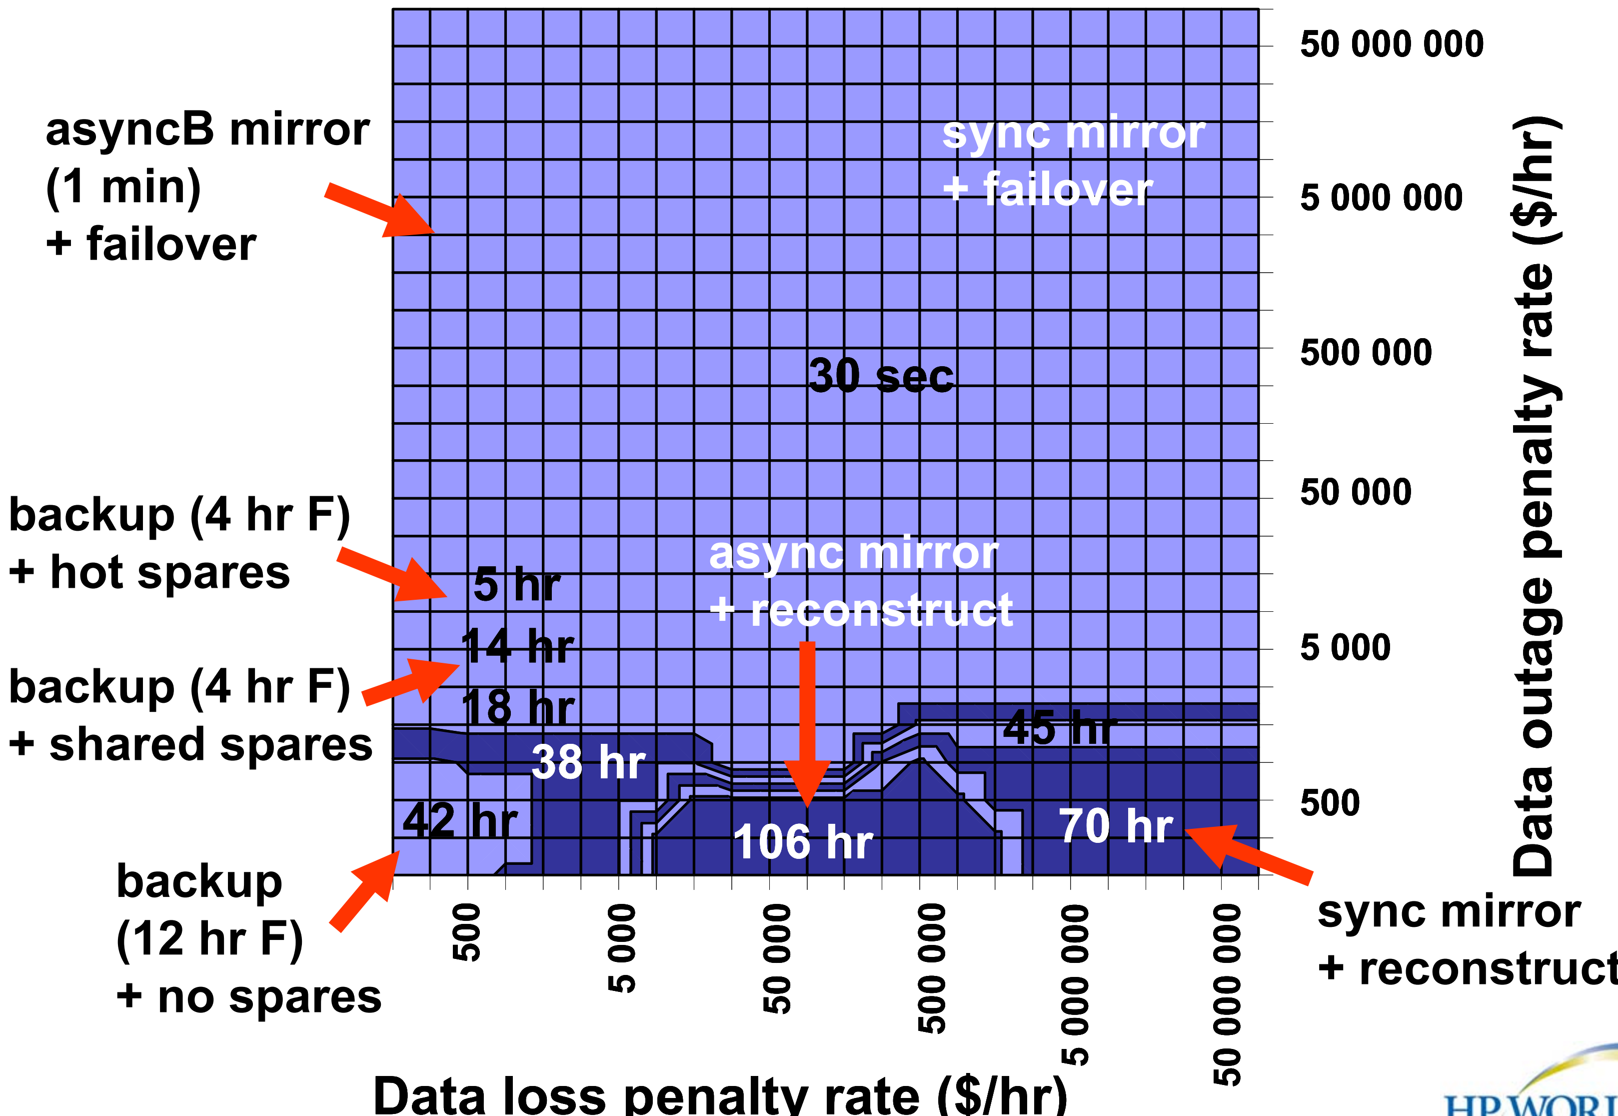

Design recent data loss

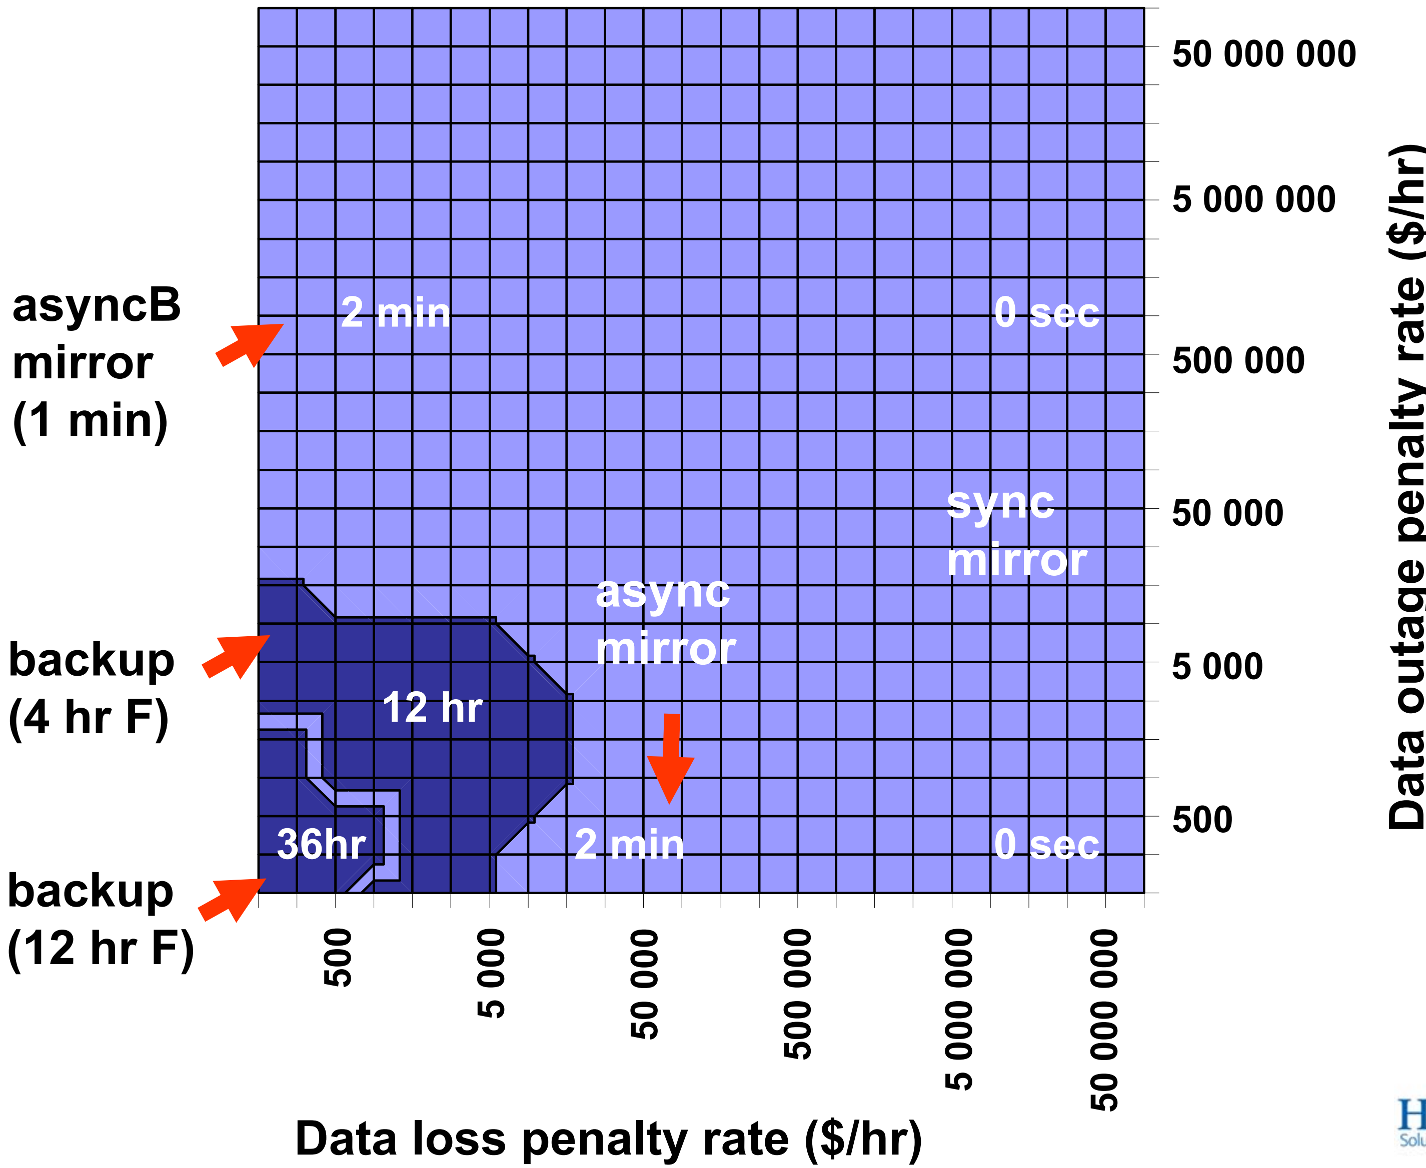

HP WORLD 2004 Solutions and Technology Conference & Expo

Design overall (outlay + penalty) costs

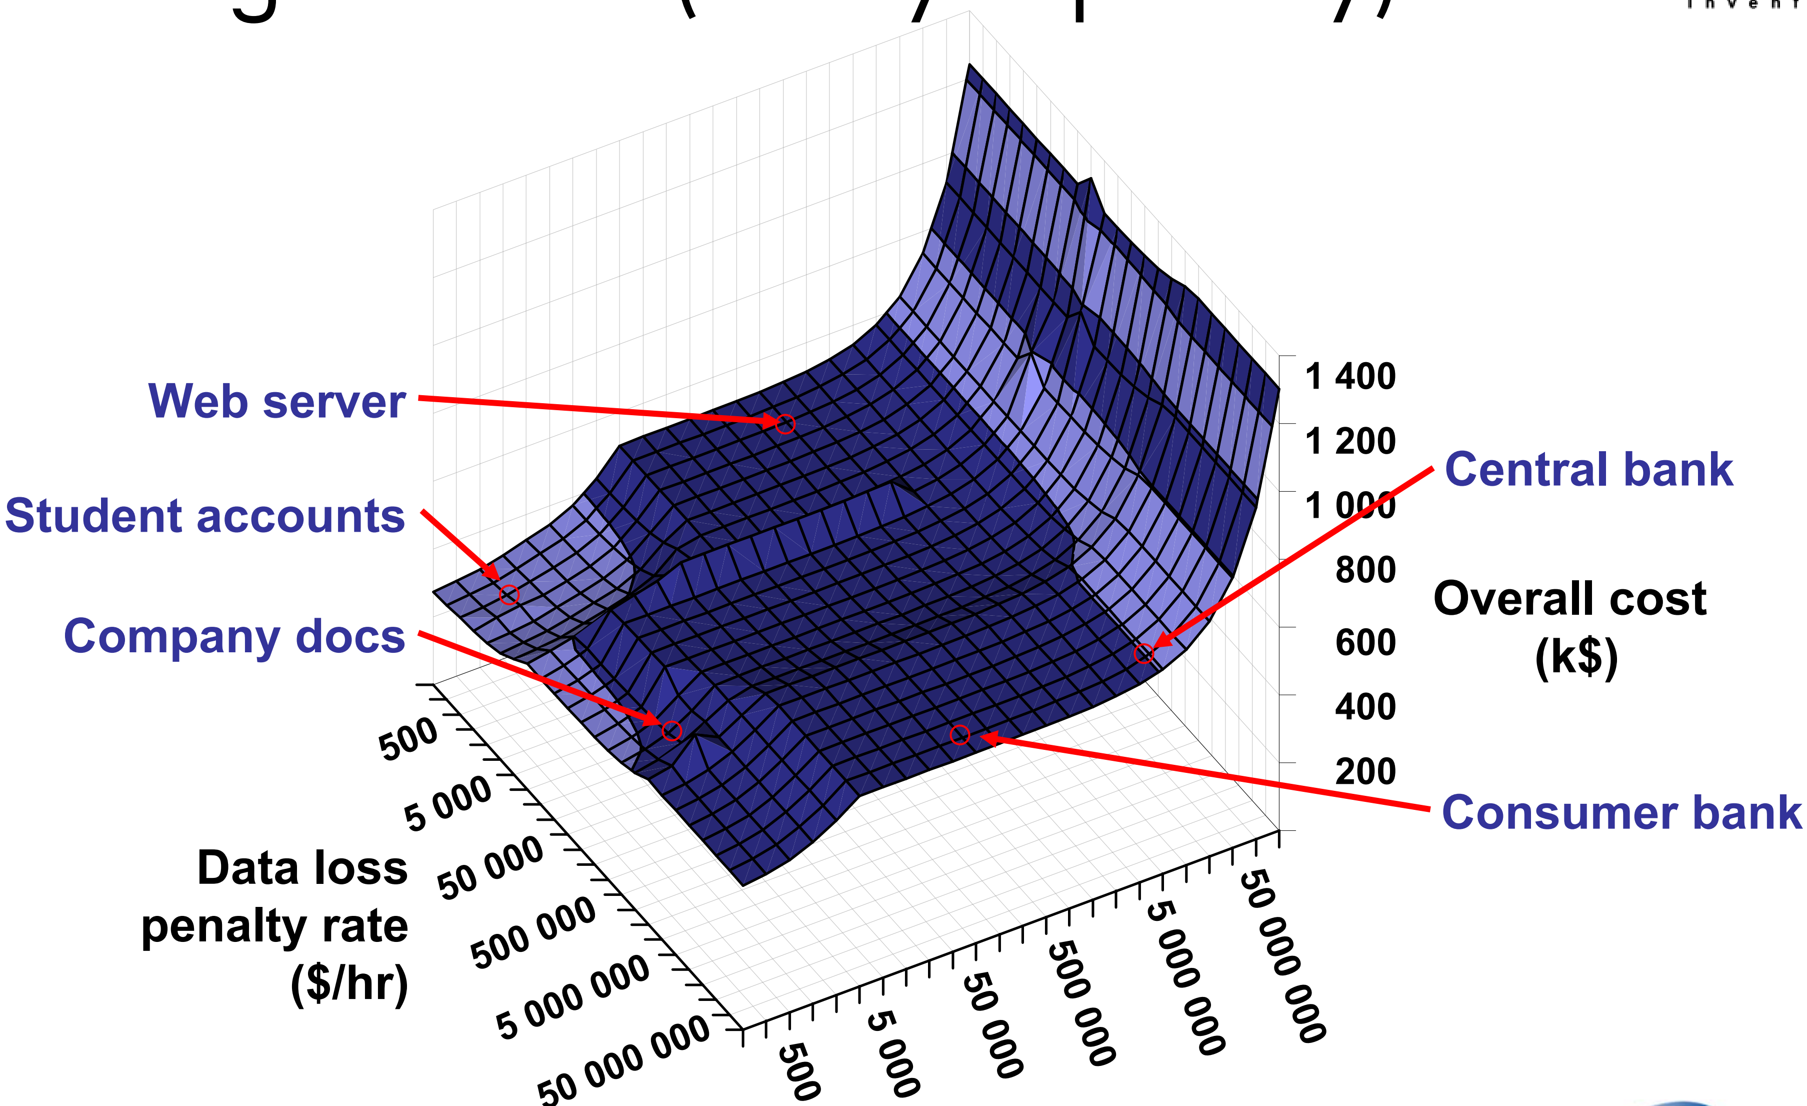

Your feedback

How do you design dependable storage systems today?

- How do you pick RTOs and RPOs?
- What other requirements do you consider?
- How do you determine how much you're willing to pay for the solution?
- How do you trade off RTO/RPO requirements and solution costs?
- How long does this process take?

How to describe business requirements?

- Some possibilities:
  - RTO and RPO
  - RTO and RPO, plus method to convert to $
  - Penalty rates
    - $ / hr downtime, $ / hr recent data loss
  - Penalty rates as a function of duration
    - Ex: 5 minutes vs. 1 hour vs. 8 hours
  - Penalty rates for degraded mode performance
    - Ex: 0%, 50%, 75% of normal performance
- Do you have other design requirements not reflected here?
  - Ex: interoperability, regulatory requirements

How much workload info possible?

- Workload characteristics
  - Capacity
  - Access rates
  - Update rates (both with and without overwrites)
  - Burstiness

- Would you be willing to run standard tools to trace and analyze workload requirements?

What info for decision-making?

- System dependability
  - Recovery time, recent data loss under failure scenarios
- Financial ramifications
  - Outlay costs
  - Penalty costs under failure scenarios
- Comparison of alternatives for a given set of requirements
- Design choice sensitivity to:
  - Business requirements
  - Workload characteristics
  - Failure scenario frequencies

How would you use a tool like this?

- Evaluation of existing designs
- What if scenario analysis
- Automated design choices
- Dependability choice exploration

Anything else you'd like to share?

Conclusions

- Automatically designing storage systems to meet dependability goals is achievable
  - Evaluate business impact of a particular solution
  - Pick best solution for specified inputs
  - Explore sensitivity of solution choice and cost to input specification
- Potential benefits for HP's customers
  - Provide ability to assess dependability of customer configurations
  - Significantly reduce time to identify appropriate solution
  - Enhance customer understanding of financial impacts of solution dependability
- Further details available:
  - http://www.hpl.hp.com/SSP
  - kimberly.keeton@hp.com

Acknowledgements

- Data dependability research is joint work with:
  - Dr. Dirk Beyer, HP Labs
  - Dr. Jeffrey Chase, Duke University
  - Dr. Cipriano Santos, HP Labs
  - Dr. Arif Merchant, HP Labs
  - Dr. John Wilkes, HP Labs

invent

HP WORLD 2004
Solutions and Technology Conference & Expo

Co-produced by:

interex
shared knowledge • shared power

encompass
AN HP USER GROUP

RECOMMENDED TRAINING VENUE FOR THE
HP Certified Professional