How to configure the links needed for long distance replication

Karen A. Fay

Senior Technical Consultant

Agenda

- Data Availability – Business Continuity
- Long distances replication solutions
- The essentials
- Tools
- Distance impact
- Configuration examples
- Tips

Data availability

Why we care…the cost of downtime

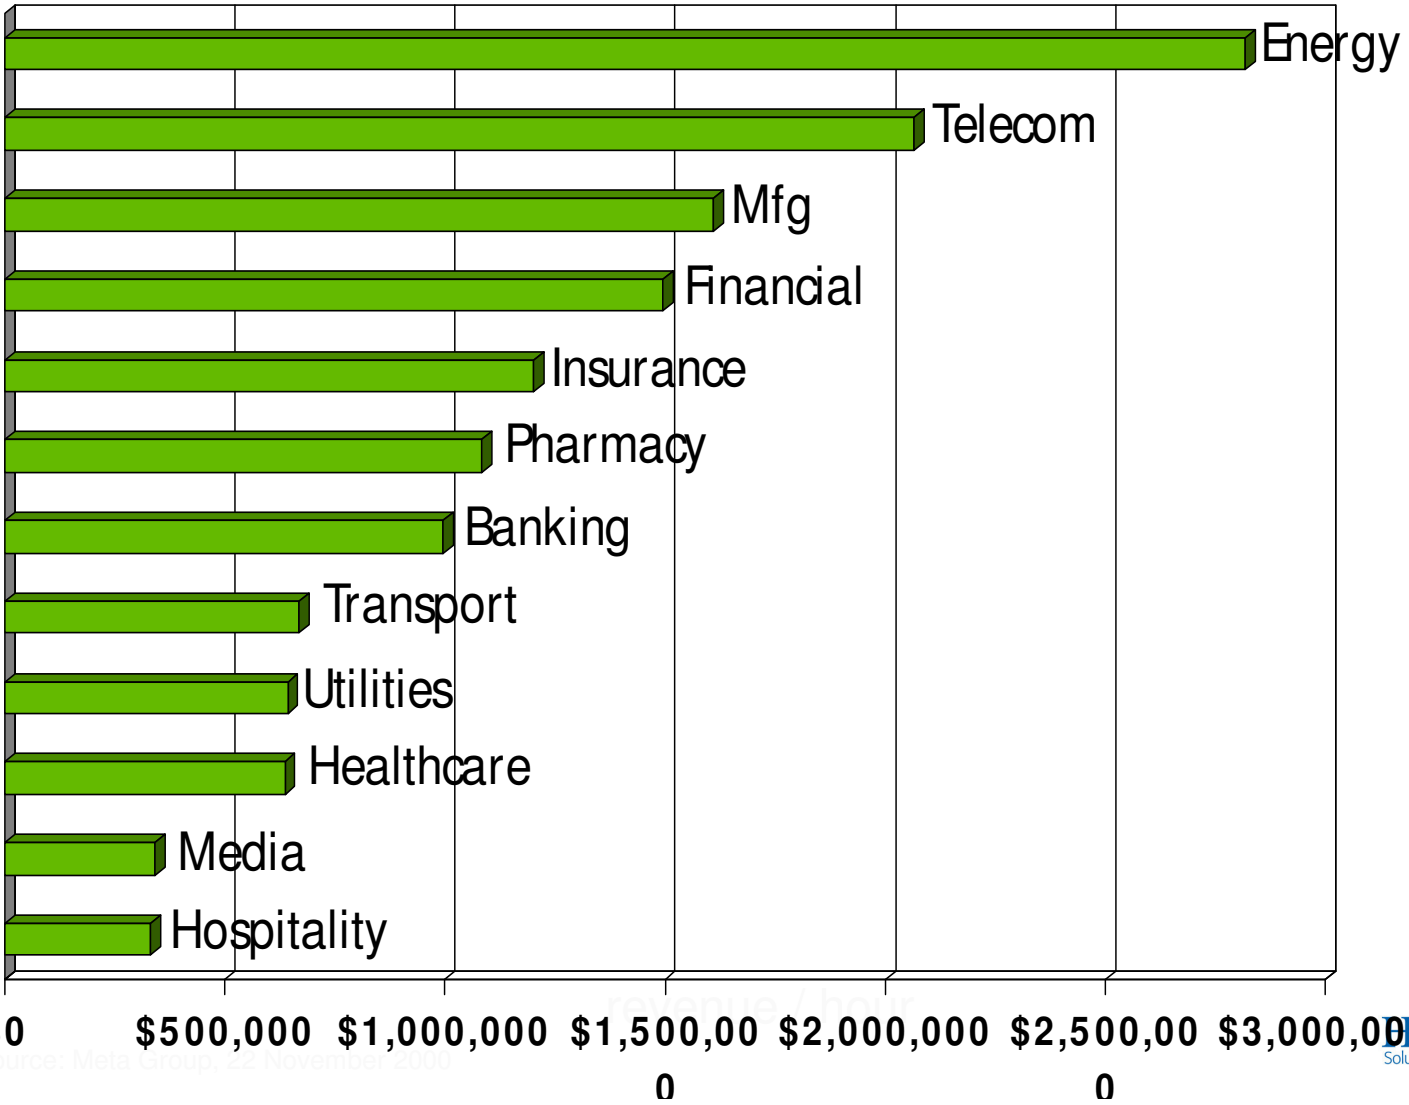

Source: Meta Group, 22 November 2000

Causes of data loss or downtime

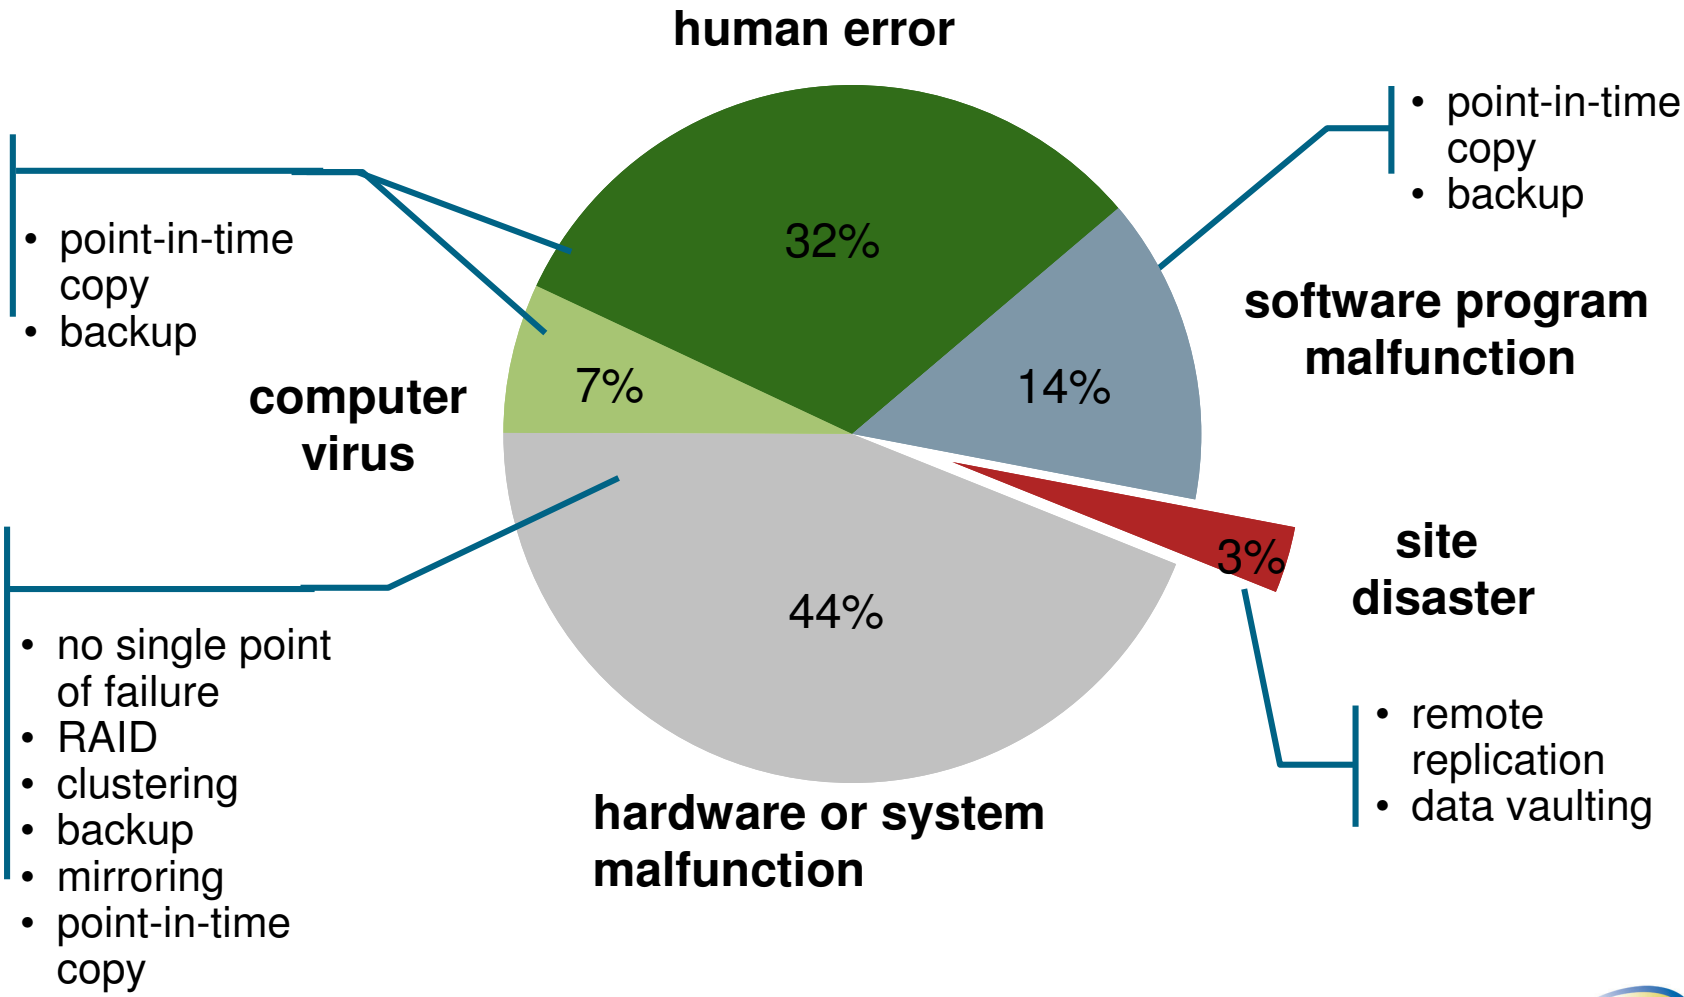

source: Ontrack, a data availability service provider

Risks to information availability

the site failure event spectrum

regional event

metropolitan area event

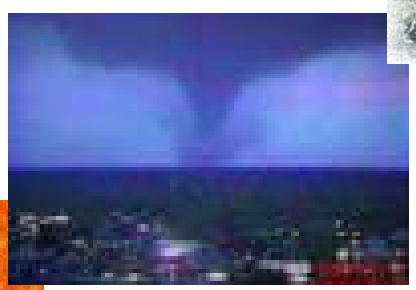

building level incident

or worse…

Choosing the right recovery strategy

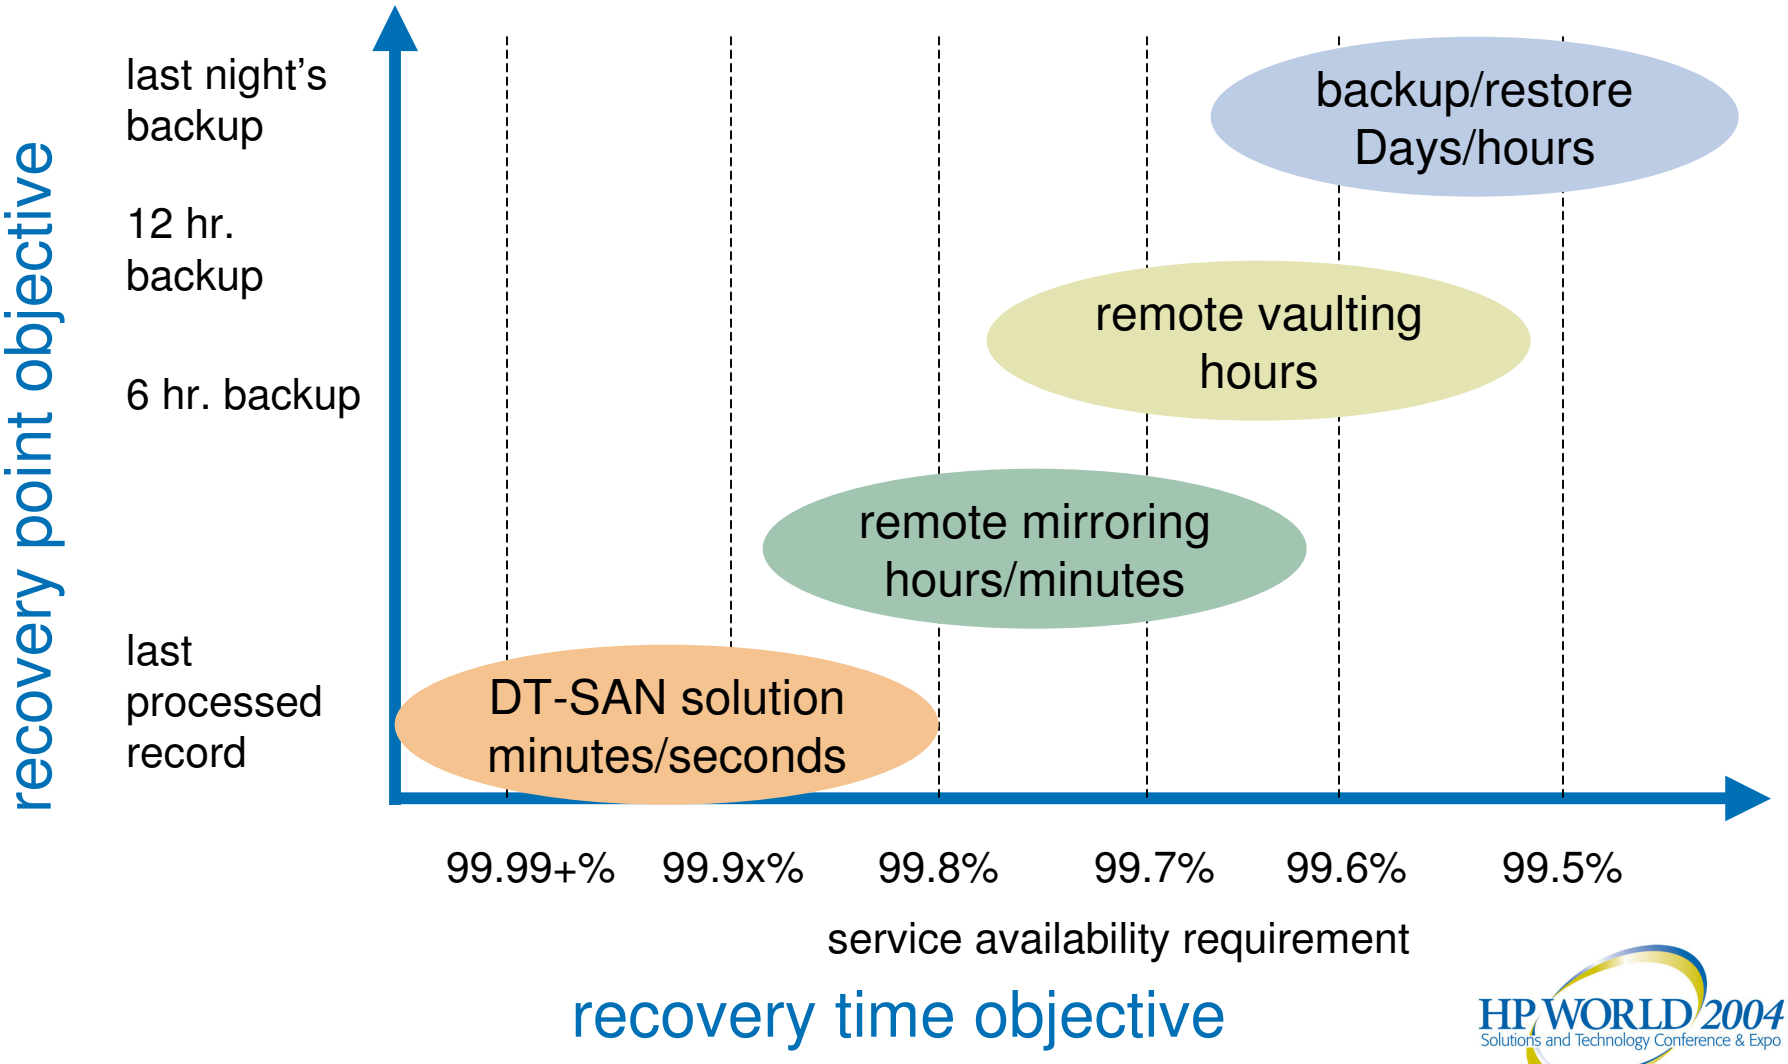

HP long distance replication solutions

Data replication approaches vary

- volume shadow technology
- OpenView Storage Mirroring

- pool and manage heterogeneous arrays
- replication across unlike arrays and hosts
- O/S dependent
- host overhead
- TCP based
- IP based clustering

- Data Replication Manager (MA/EMA family)
-Continuous Access XP
- Continuous Access EVA

- logically pool array
- dynamic LUN optimization
- like for like arrays
- vendor specific, no interoperability

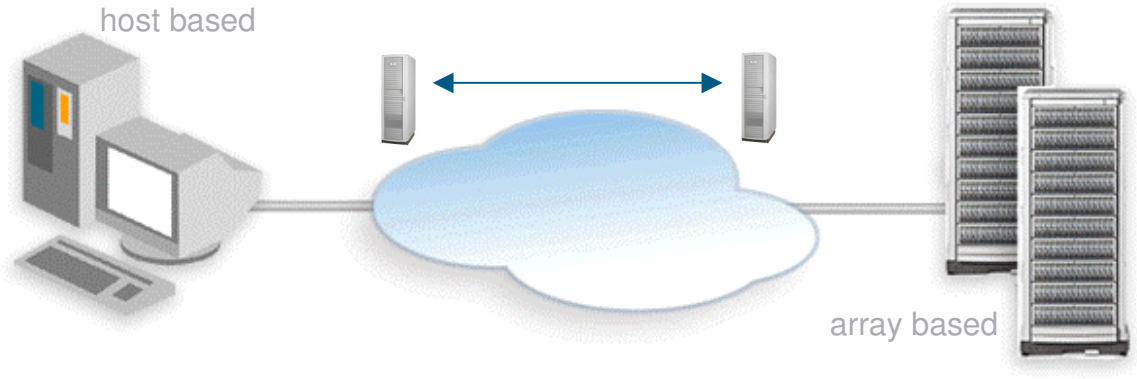

The essentials

Ground rules

- 40-45% utilization of links
- Equal # of links for each fabric
- Size for daily activity as well as normalization
- Know what you are trying to protect against:
  - fire
  - flood
  - earthquake
  - or worse
- Distance is the enemy in reducing latency
- Replication is done on a volume (LUN) basis

Basic elements to know

- Distance between sites
- IOPS
- Average I/O size
- Link speed
- # of links

Tools

Tools

- Distance calculators
  - Ping to the other site
  - Qwest.com for link latency/distance
  - Point-to-point
    - Driving distance * 1.5
  - Routed networks
    - Driving distance * 2.25

Tools

- IOPS & average I/O size – writes only
  - OVSAM performance optimizer
  - Operating system tools
    - IOSTAT
    - PERFMON
    - Etc…

Tools

- Solution specific tools
  - DRM Performance Estimator
  - CA EVA Performance Estimator
    - http://h18006.www1.hp.com/products/storage/software/conaccesseva/relatedinfo.html
  - Design guides
    - http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?contentType=SupportManual&locale=en_US&docIndexId=179911&taskId=101&prodTypeId=12169&prodSeriesId=316118

Distance impact

Impact of distance

- Speed of light is 5 microseconds per kilometer one way
- SCSI write with EVA replication is one round trip
  - *Here's the data*
  - *Got the data*
- Read or write over 100km adds 2 milliseconds
  - 15 k RPM drive rotational latency is 2 mSec
  - 7200 RPM drive rotational latency is 4 mSec
  - 100 km reduces performance similar to slower drive

Configuration examples

Sample configuration

- Distance: 200 KM (125miles)
- IOPS: 250 writes (total for all Copy Sets - volumes)
- Ave I/O size: 8 KB
- # Remote Copy Sets: 12
- Ave LUN size: 250 GB
- Total capacity: 3 TB
- Running Continuous Access on EVA

Requirements for daily activity

- **2 1GE-IP/1000 links:**
  - 3% utilized/link
  - 375 IOPS/link
- **2 OC3/100 Mb links:**
  - 24% utilized/link
  - 296 IOPS/link

2 T3 IP/44 Mb links:

42% utilized/link

228 IOPS/link

6 10 Mb IP links (to get utilization down):

68% utilized/link

84 IOPS/link

Don't even think about T1 links, since…

T3 = 44 Mbs, T1 = 1.54 Mbs

Requirements for normalization stage

- **2 GE-IP links:**
  - 28% utilized/link for 1 stream-Copy Set, so 3 Copy Sets/link at a time
  - 6 Copy Sets in parallel/across 2 links = 2 phases
  - 101 GB/hour/Copy Set, so…2.5 hours for each phase
  - Total: 5 hours
- **2 OC3/100 Mbs links:**
  - 77% utilized, so…1 Copy Set/link at a time
  - 2 Copy Sets in parallel/across 2 links = 6 phases
  - 28 GB/hour/Copy Set, so…9 hours for each phase
  - 6 phases*9 hours
  - Total: 54 hours
  - or 4 OC3 links = 27 hours
- **2 T3/44 Mbs links:**
  - 88% utilized, so…1 Copy Set/link at a time
  - 2 Copy Sets in parallel/across 2 links = 6 phases
  - 14 GB/hour/Copy Set, so…18 hours for each phase
  - 6 phases * 18 hours
  - Total: 108 hours
  - or 12 T3/44 Mbs links = 18 hours

Tips

Misc tips

- Qwest.com
- Price link from both ends
- Get SLA for needed bandwidth
- Normalize
  - on same site then apply write log
  - restore backup then apply write log
- Calculate each side for bi-directional configurations and add together
- Compression should not be considered during configuration of links

Don't let this happen to you…

Continuous Access is the answer


invent