Session 3709 Designing business copy solutions for the EVA

Don Fraser
Disaster Tolerant Solutions & Test Architect
NSS CA EVA Systems Integration & Test
Hewlett-Packard

Agenda

- Part one – a quick review
  - Overview of how snapclones and snapshots work
  - Basic versus Enhanced BC support
  - Supported configurations
- Part two – elements of design
  - Understanding the requirements
  - Understanding the limitations

Terms and concepts

- Snapclone
  - A point in time physical copy of the data
  - Only one active at a time

- Snapshot
  - Both Fully Allocated and Space Efficient
  - A point in time ‘virtual copy’ of the data
  - Up to 7 per source Vdisk

Terms and concepts (cont)

- Crash Consistency
  - Point in time view of data as if server just crashed
    - Data in server cache is lost
  - In-complete transactions expected
    - Application must be able to clean up
- Transaction Consistency
  - Quiesce or shutdown Application
  - Flush server cache(s), and array cache
  - No incomplete transactions expected
- Currency
  - Up to date

Terms and concepts (cont)

- Two interfaces:
  - Command View EVA
    - Basic Support
    - No host agent
    - Optional SSSU interface
  - Business Copy for EVA
    - Enhanced Support
    - Host Agent
    - Scheduler
    - Allows coordination of Application and BC operation

New options with VCS 3.02

- Cross VRAID Snapshot
  - Must be in within same diskgroup as source
  - Choose Fully allocated or space efficient
  - Select VRAID type of Snapshot regardless of parent
    - VRAID 5
    - VRAID 1
    - VRAID 0
- Cross VRAID Snapclone
  - Choose Destination Disk group
    - Online or near-online
  - Choose VRAID type of snapclone regardless of parent
    - VRAID 5
    - VRAID 1
    - VRAID 0

Snapclone Details

- Physical copy of a production volume.
- Process:
  - Quiesce the application
    - Need to quiesce application to insure all pending transactions are written to disk.
    - For example “File consistent” or “Transaction consistent”
    - Otherwise get “Crash consistent” (e.g. server crashed)
  - Create clone,
  - Resume application
  - Present to a separate server
- The clone will contain exact point in time copy of production data. The clone can be used to run tasks in parallel with production data but on a separate server
- Data copy is via the EVA backend, not SAN
  - Other activity may limit bandwidth available for copy

Creating a Snapclone

- Create Snapclone
  - data copies in background

  => a post create point-in-time normalized copy of source

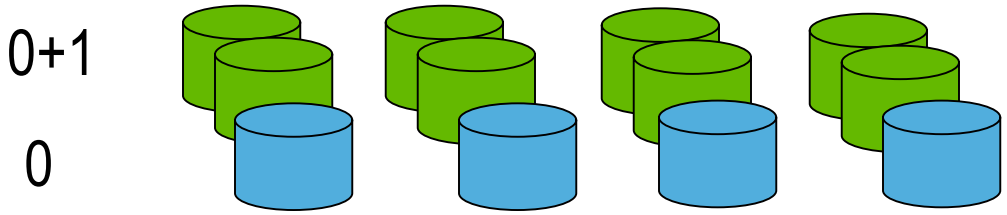

VSnap Details

- Virtual copy of production volume - no data moves until production volume changes
- Quiesce application before Snap
- Need to quiesce application to insure all pending transactions are written to disk.
  - For example “File consistent” or “Transaction consistent”
  - Otherwise get “Crash consistent” (e.g. server crashed)
- Contains exact point in time copy of production data
- Can Snap VRAID 0, VRAID 1, or VRAID 5 Vdisks.
- The Production volume and the Snapshot are Read/Write.
- The Production volume will always contain latest data.
- A temporary volume will contain changes made to the Snapshot volume.
- Data movement is via the EVA backplane

How do Snapshots Work?

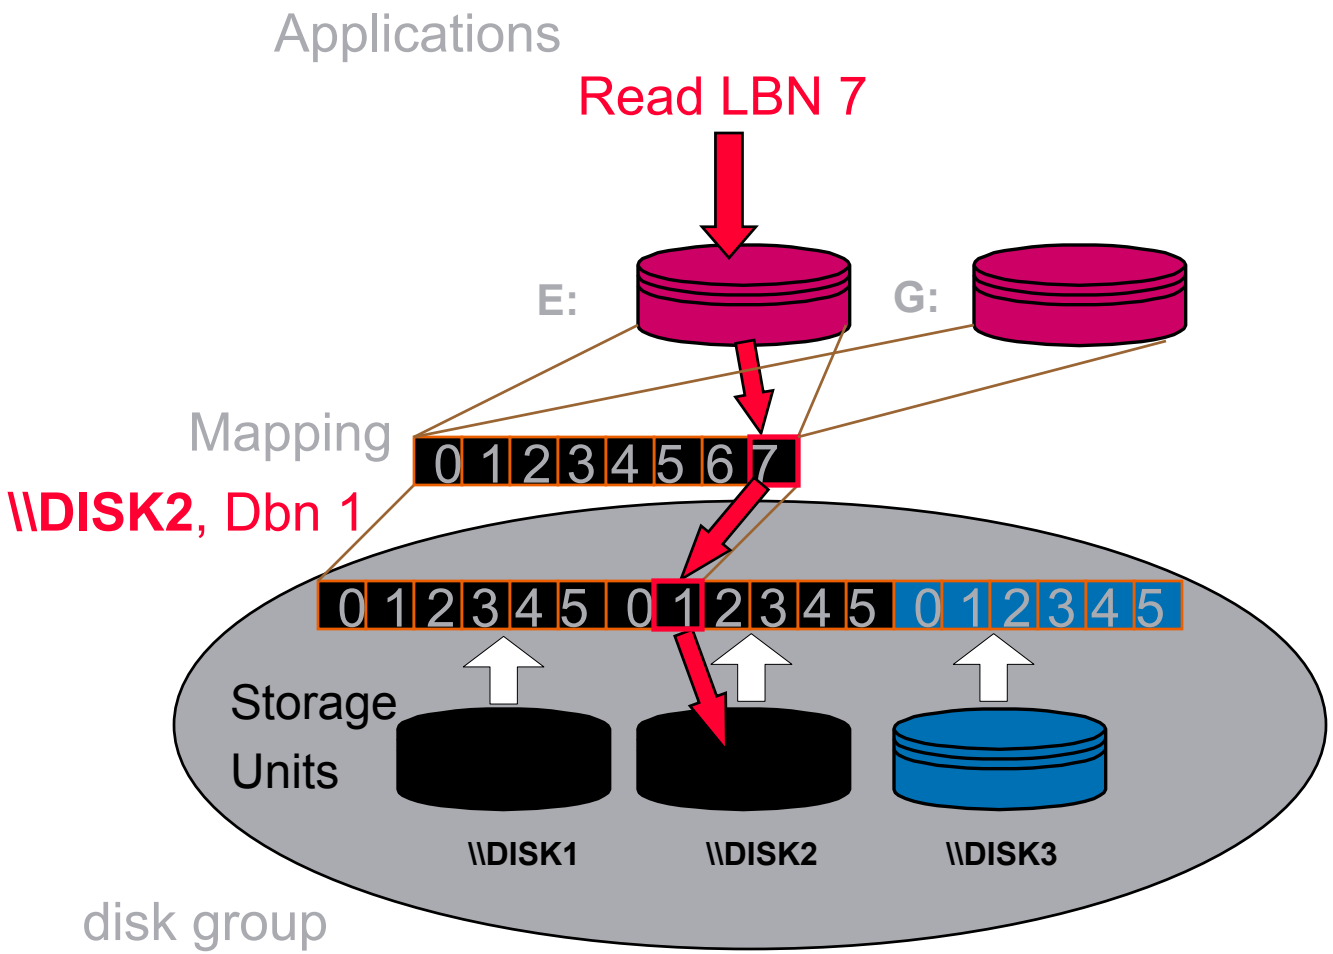

How do Snapshots Work?

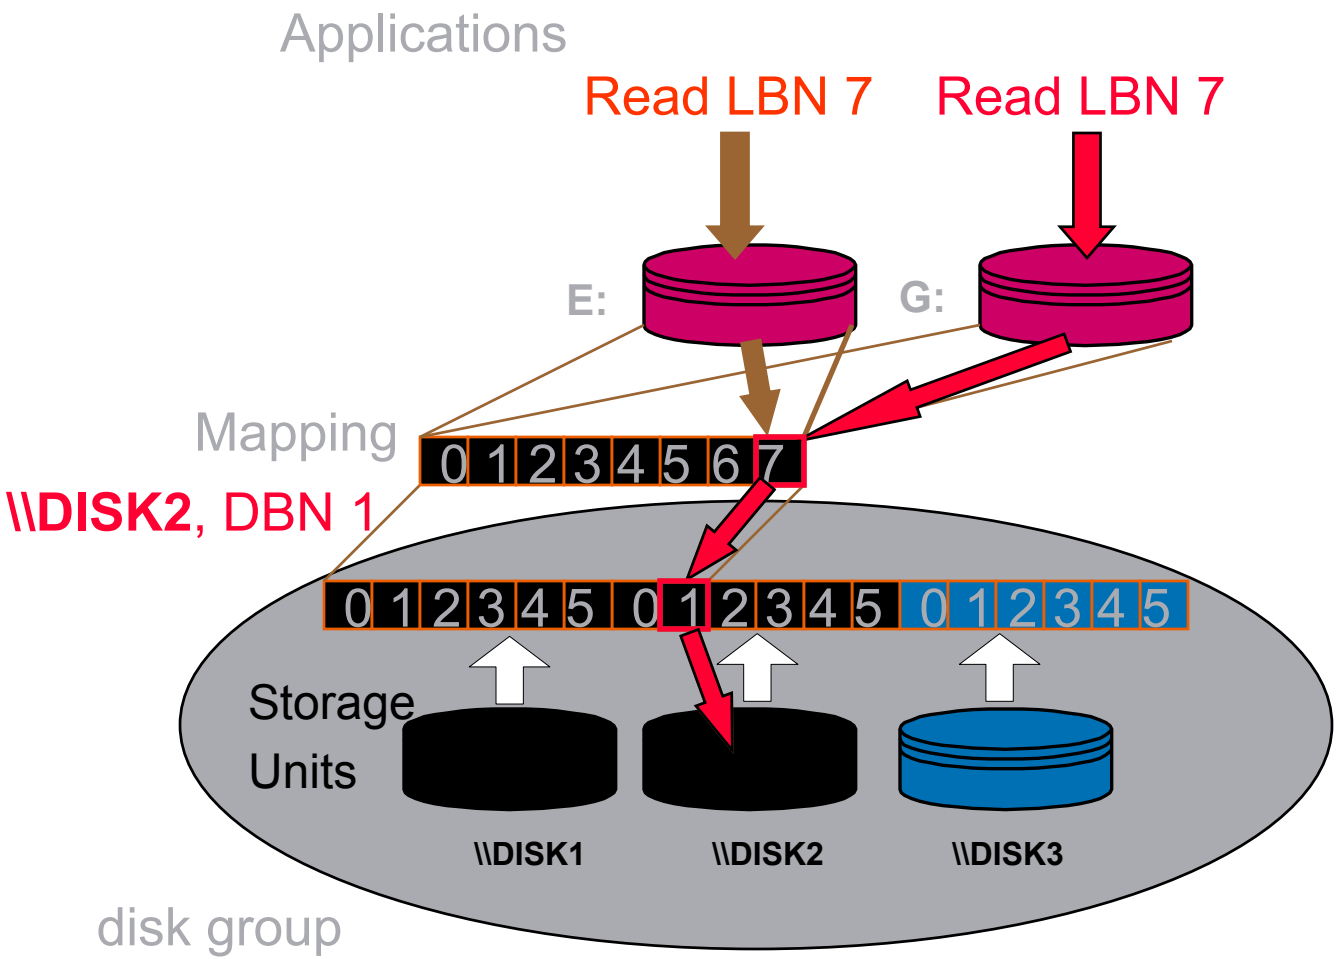

How do Snapshots Work?

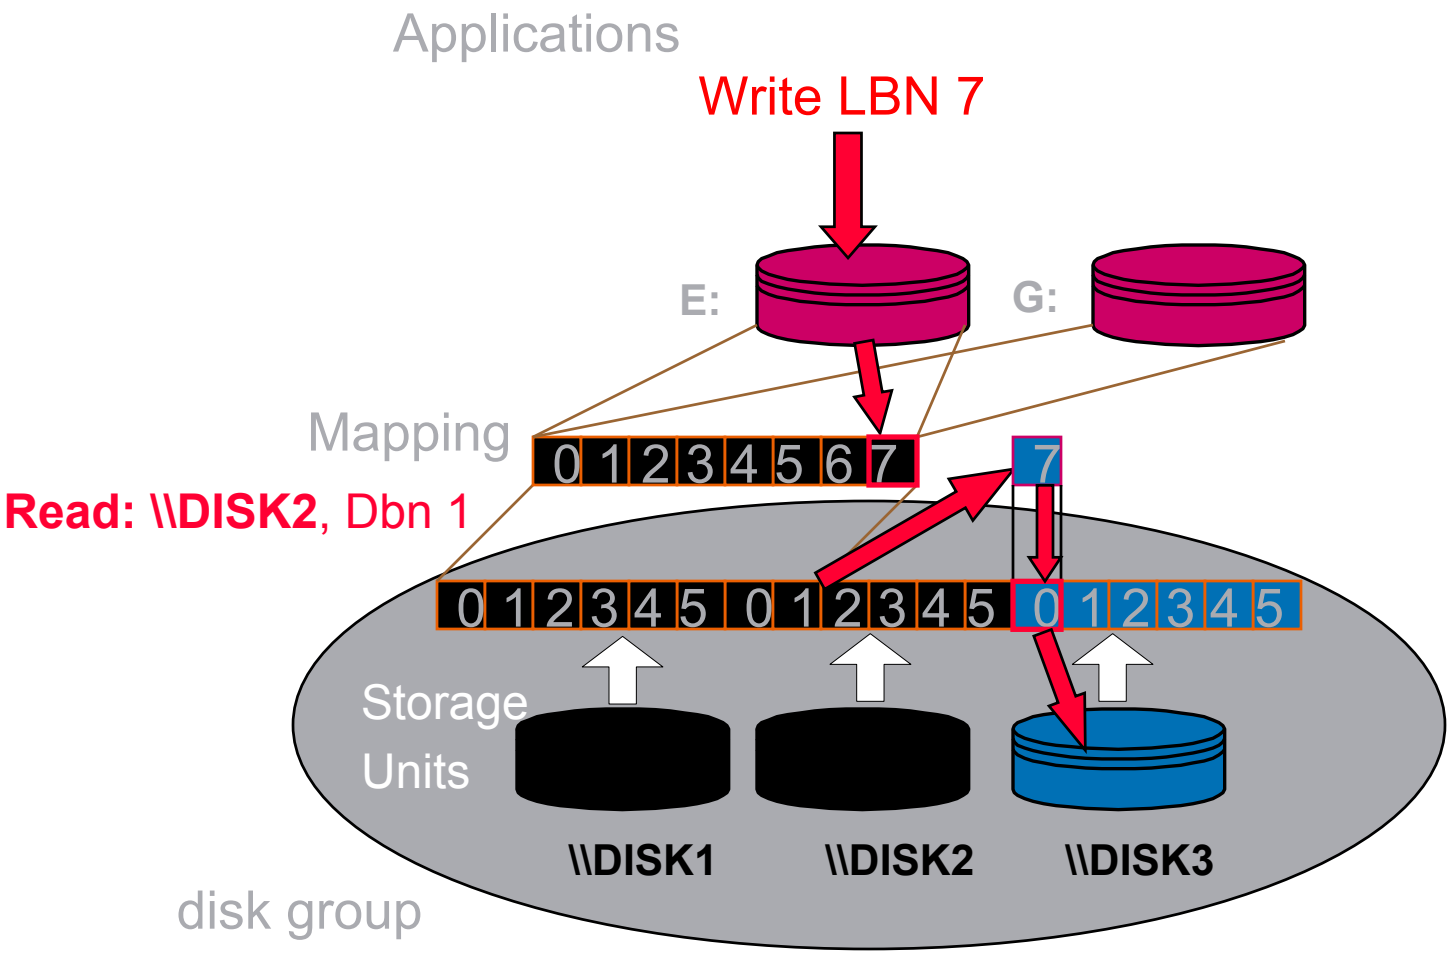

How do Snapshots Work?

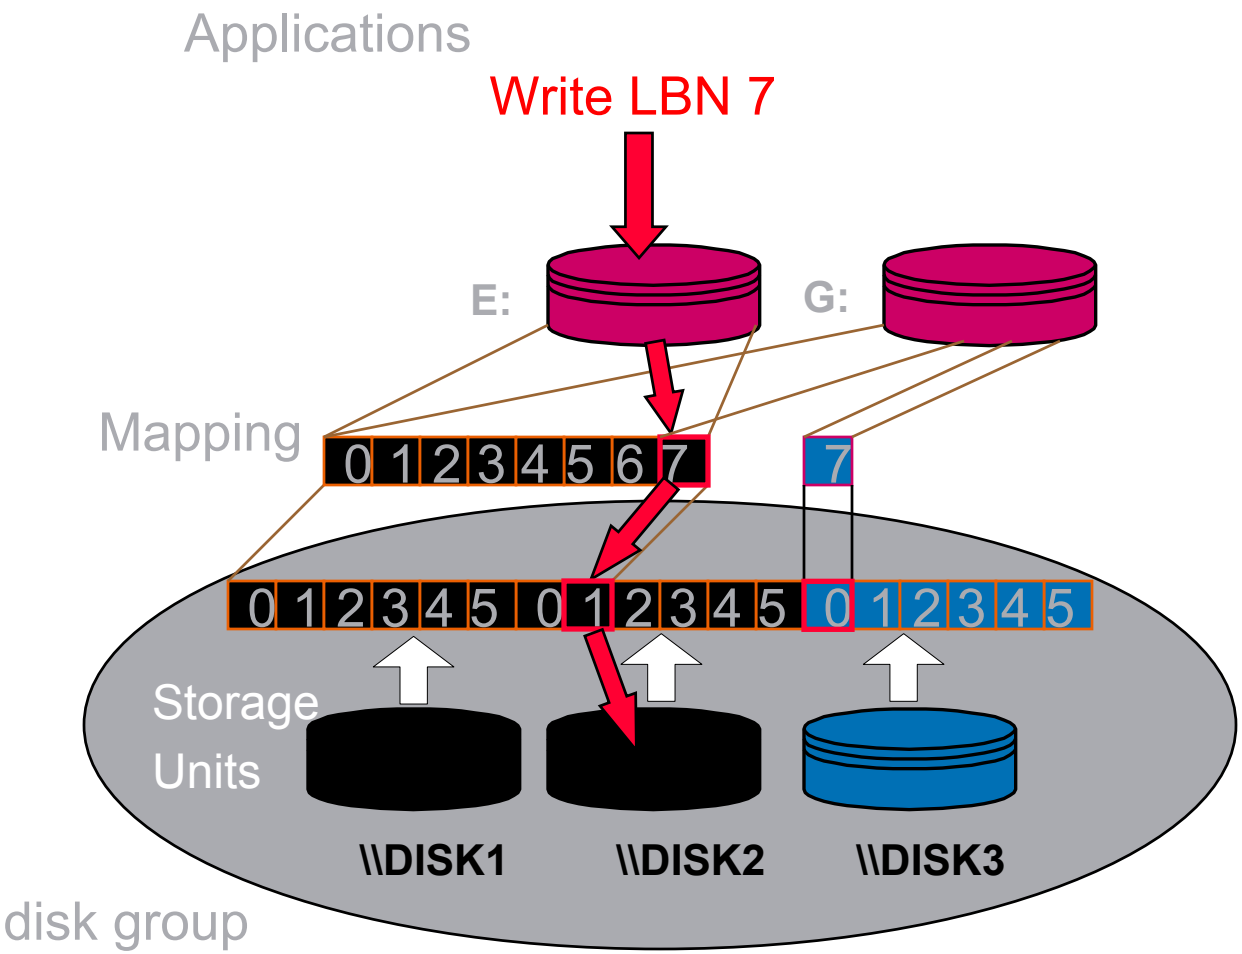

How do Snapshots Work?

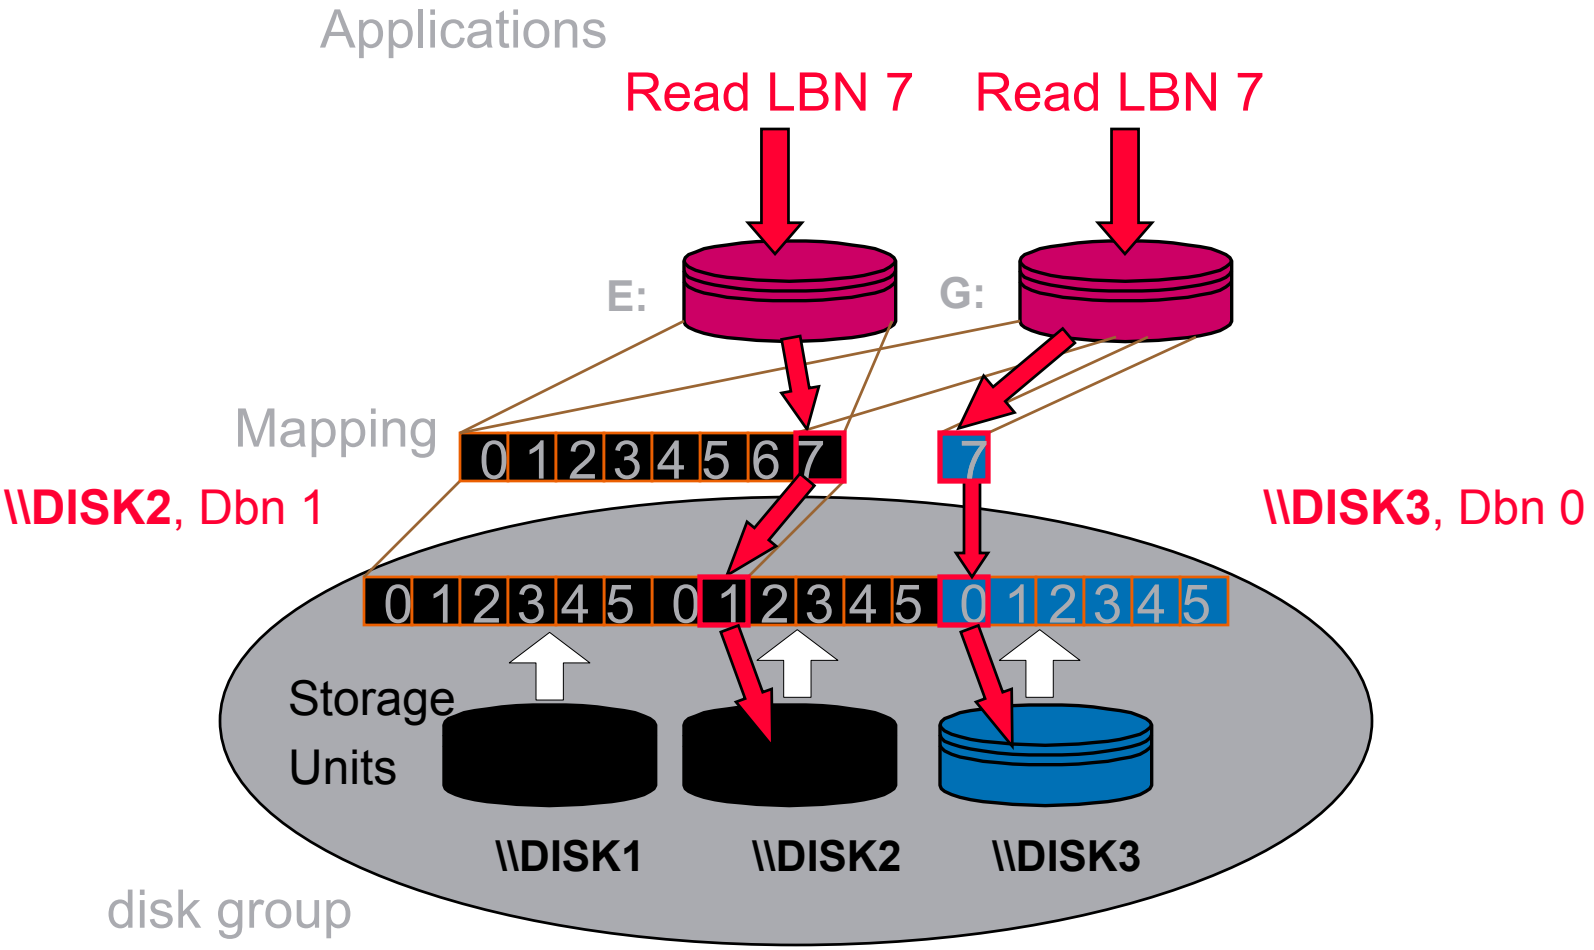

Agenda

- Part one – a quick review
  - Overview of how snapclones and snapshots work
  - Basic versus Enhanced BC support
  - Supported configurations
- Part two – elements of design
  - Understanding the requirements
  - Understanding the limitations

Basic support

- Requires BC license for EVA
- Available with Command View EVA UI
- No host assists
  - No host agent
  - No job engine or scheduler
- Might use SSSU

Basic support is available for

- HP HP-UX 11.0, 11i, 11i v2
- HP OpenVMS 7.2-2, 7.3-1, 7.3-2
- HP Tru64 5.1a. 5.1b
- IBM AIX 4.3.3, 5.1, 5.2
- Microsoft Windows: NT; 2000; 2003 (32 & 64 bit)
- Novell Netware 5.1, 6.0, 6.5
- Red Hat AS 2.1 (32 & 64 bit)
- SUN Solaris 2.6, 7, 8, 9
- SuSE SLES 7, 8 (32 bit & 64 bit)
- United Linux V1.0 (32 bit)

Enhanced support

- Starts with Basic support

- Adds
  - BC specific user interface
  - Host Agents
  - Job scheduler

Enhanced support is available for

- HP HP-UX 11.0, 11i, 11i v2
- HP Tru64 5.1b
- IBM AIX 5.1, 5.2
- Microsoft Windows: NT; 2000; 2003 (32 & 64 bit)
- Red Hat AS 2.1 (32 & 64 bit)
- SUN Solaris 7, 8, 9
- SuSE SLES 8 (32 bit & 64 bit)
- United Linux V1.0 (32 bit)

Agenda

- Part one – a quick review
  - Overview of how snapclones and snapshots work
  - Basic versus Enhanced BC support
  - Supported configurations
- Part two – elements of design
  - Understanding the requirements
  - Understanding the limitations

Why make business copies?

- For online backup?
  - Use physical copy, Snapclone
- For file recovery?
  - Use virtual copy, fully allocated Snapshot
- As source for backup to tape?
  - Use virtual copy, space efficient Snapshot

- Crash Consistency or Transaction Consistency?
  - Time to quiesce or shutdown application.
  - Time to re-start application.

Understand the limitations

- Performance limits of
  - The controller
  - The backend
  - The diskgroup
    - Source
    - Destination
- What else is happening?
  - Application I/O
  - CA
  - Leveling
  - Rebuild

Performance limits

- Sequential access
  - Limited by bandwidth of individual disk and degree of parallelization
    - Use 1000 requests per second
    - Regardless of size
    - Regardless if read or write
- Random access
  - Usually limited by number of drives in disk group
    - Use 150 requests per second * number of drives in disk group
    - Regardless of size
    - Regardless if read or write

Examples

- Disk group of 10, 10 K RPM 146 GB Disks
  - Sequential access:
    - 1000 requests per second
  - Random access:
    - 150 requests per second * 10
    - Therefore disk group is capable of 1,500 requests per second
- Disk group of 100, 10 K RPM 146 GB Disks
  - Sequential access:
    - 1000 requests per second
  - Random access:
    - 150 requests per second * 100
    - Therefore disk group is capable of 15,000 requests per second

Process takes time

- EVA 5000 Array with 3 disk groups
  - #1 – 24 250 GB near-online disk drives
  - #2 – 24 300 GB online 10K RPM disk drives
  - #3 – 36 36 GB online 10K RPM disk drives
- Each diskgroup contains 2 Vdisks
  - 250 GB Vraid 1
  - 250 GB Vraid 5
- For the VRaid 1 in disk group #1, measure how long it takes to create snapclone
  - Near-online to near-online (same diskgroup)
  - Near-online to online (different diskgroup)

Process takes time -- cont

- With nothing else happening during the test
  - No host I/O
  - No leveling
- Time to create
  - Near-online -> Near-online => 67 minutes
  - Near-online -> On-line => 61 minutes
- Time to delete
  - Near-online => 39 minutes
  - On-line => 36 minutes

Where to find documentation and software

- The StorageWorks Business Copy (BC EVA) web page:

  http://h18006.www1.hp.com/products/storage/software/bizcopyeva/index.html

Then click on technical documentation to find additional resources

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