# SNIA VDBENCH Emerald Overview

Steven A. Johnson

SNIA Emerald™ Training

*SNIA Emerald Power Efficiency Measurement Specification*, for use in EPA ENERGY STAR®

July 14-17, 2014

SNIA™ Green Storage Initiative

SNIA Emerald™

# Agenda

- Introduction to SNIA Emerald
- Phases test
  - Pre-fill
  - Warm up
  - Optionally time for tiering to move things around
  - Hot band test
  - 4 corners test
    - Small block random read and write
    - Large block sequential read and write

# Determining target storage

- First part of the process is the configure your storage.
- Determine your optimal configuration (Raid 1, Raid 5, ??, Stripe size, Volume Manager settings, etc)
- Determine the amount of total capacity
- Export 56+% of the storage or all of it and let vdbench execute against 56+% of capacity (sd parameter “range=(0,57)”
- List all the luns for the test and put it in a file.  Use format or Disk Manager (OS Dependent) to identify the LUNs

# Filling storage

- For the 56+% capacity, all of it needs to be initialized by vdbench
- Vdbench will create a near random pattern that has a characteristic of 2:1 compression
- Have vdbench write from beginning to end to have a known pattern across all active blocks of storage (part of the example script provided)
- No power measurements or performance will be reported on this phase

# Warm-up storage 12+ hours

- The warm-up phase designed to have the following effect on storage:
  - Get the system to a relatively steady state.  No out of the box performance for any component.
  - Cause some amount of fragmentation in systems that do “copy on write”
  - Let intelligent tiering systems study access pattern for moving blocks around
  - Optionally slow period that allows tiering systems have cycles to move data around
- Hopefully system will approximate a customer system after weeks of activity

# Hot Band run

- Hot Band workload
  - Hot band is a cache friendly workload
  - Has very hot spots that tend to stay in cache
  - More cache the more cache hit the storage should produce
  - 8 sequential streams, 5 read, 3 write make up 40% of workload
- Intended to simulate a transactional workload

# Windows Disk Manager

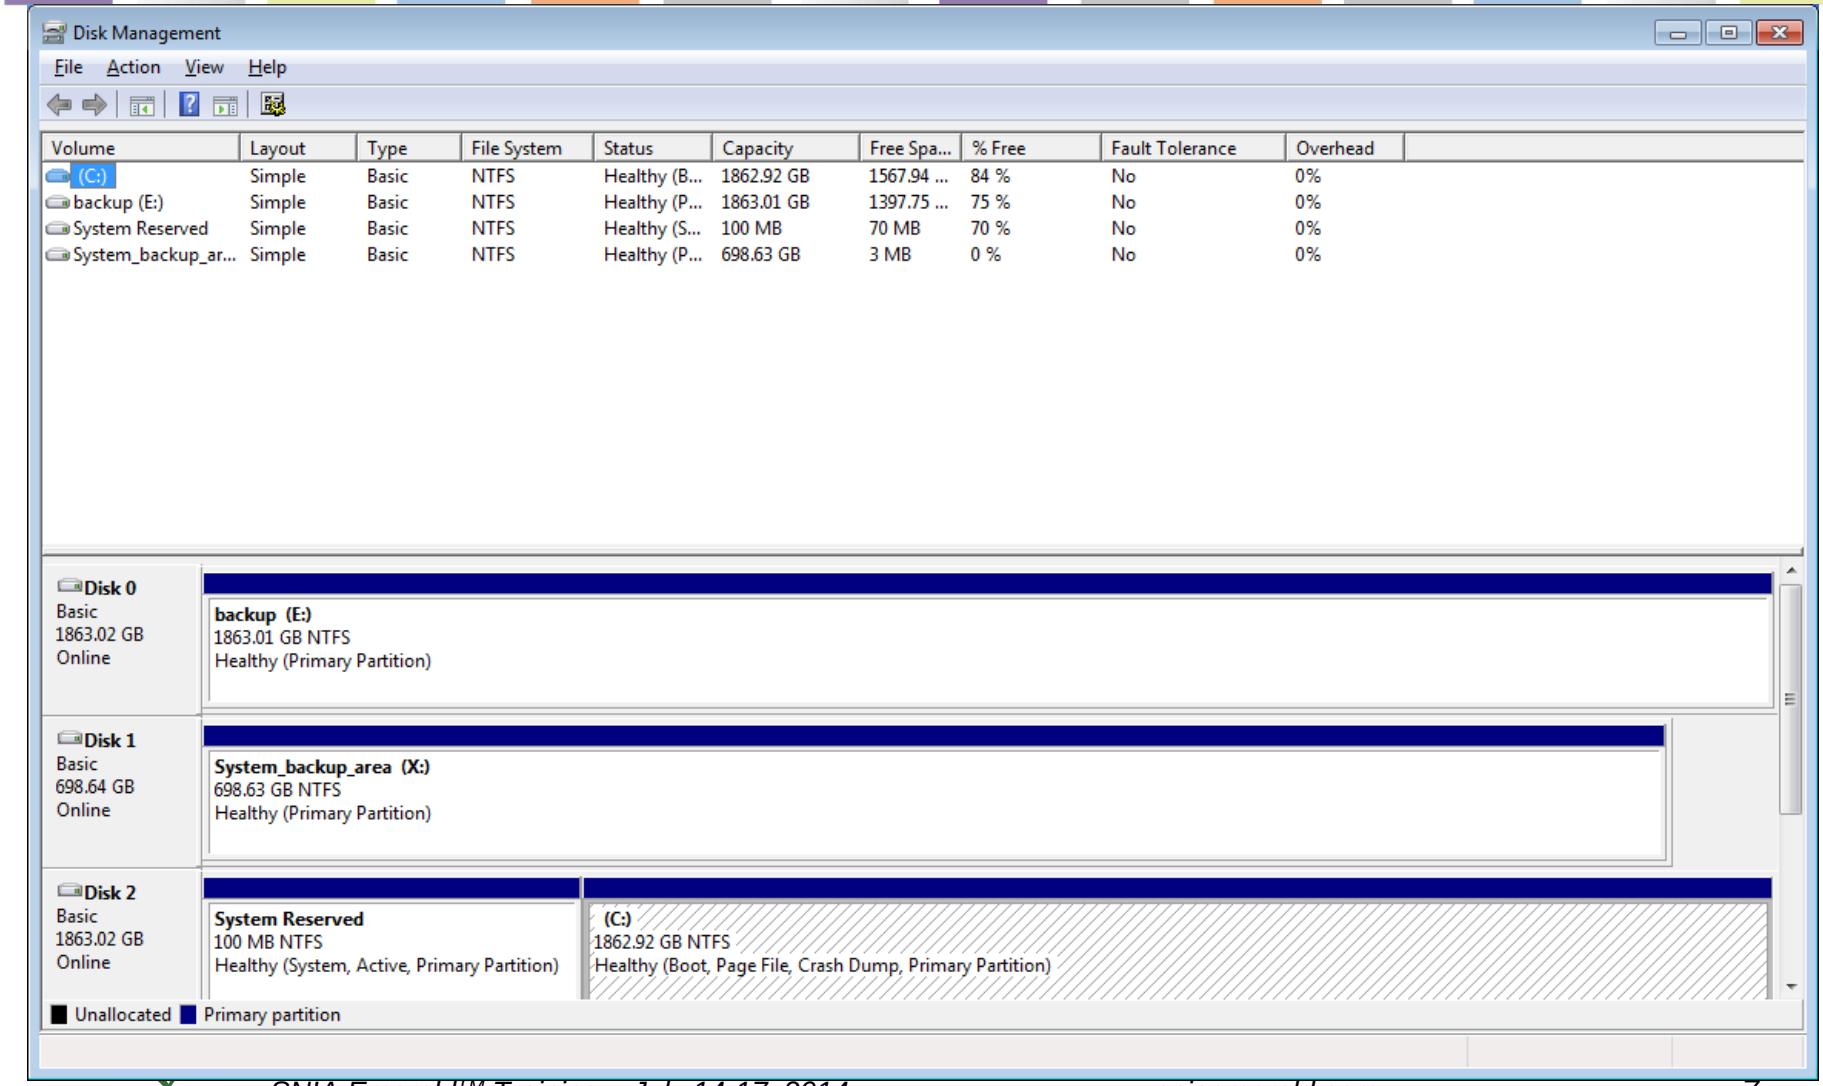

# Solaris Format

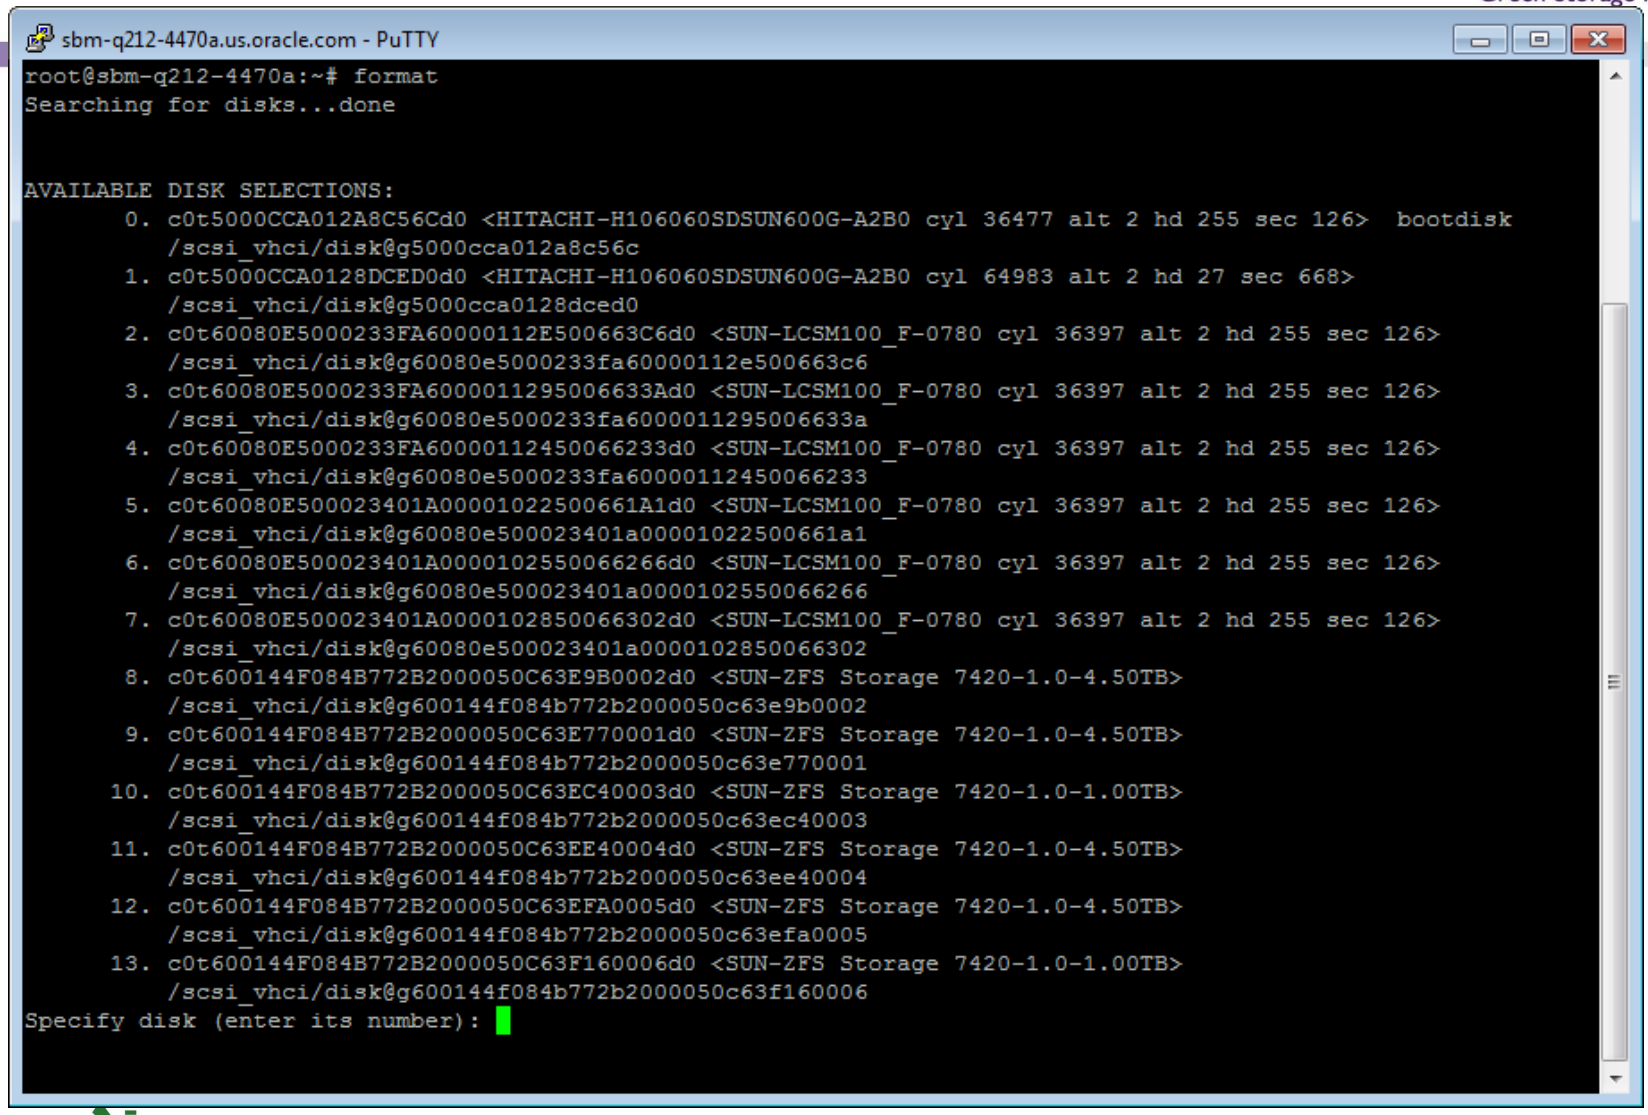

# Red Hat parted

- Linux will have various ways to find the partition table
- Need to figure out your version
- Red Hat has the “parted” command and print will display the drives and partitions on each drive
- See details
  https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/s1-disk-storage-parted.html#s2-disk-storage-parted-view-part-table

# 4 Corners 20 minutes

- Tests the extremes of workloads
  - Random Read and Write 8K transfer
  - Sequential Read and Write 256K transfer

# Phase time line

**Typical**

| Phase | Fill | Hot Band Warm-up | 4 Corners |
|---|---|---|---|
| Duration | MB capacity/MB rate/3600 | ~12Hr | 80 Minutes |

**Tiering**

| Phase | Fill | Hot Band Warmup/learning | Hot Band Semi-idle Move blocks | Hot Band | 4 Corners |
|---|---|---|---|---|---|
| Duration | | ~12Hr | TBD | TBD | 80 Minutes |

# Questions