ENERGY STAR®
Data Center Storage
Version 1.0

Stakeholder Meeting
San Jose, CA

20 July 2010

EPA

Agenda

- Draft 1 Comments Review
  - Definitions
  - Data Measurement & Output
  - Efficiency Criteria
  - Scope
  - Taxonomy
- Product Definition Proposal
- Additional Data Collection

Selected Stakeholder Comments on Draft 1

Definitions

- Low power modes may be of little relevance in some taxonomy categories if latency and response time impacts are substantial.
- For consistency, the “ready idle” definition should require that all devices be in a running, ready-idle condition, with no devices in a low-power state with reduced response time.
- The term “deep idle” does not seem to be widely used in the industry – suggest a more common term.

Data Measurement & Output

- Proposed data acquisition sampling & reporting rates are unreasonable
  - *Sampling frequency intended to facilitate data availability, not actual reporting. Consistent with server requirements.*
- Proposed measurement accuracy requirements for temperature & power are unreasonable.
- Use of off-the-shelf components would simplify integration.
  - *Intelligent PDU may address some concerns.*

Energy Efficiency Criteria

- It is important to ensure that “capacity reduction” techniques actually deliver a net energy savings.
- Power management techniques are more relevant to servers than to storage.
- Recommend starting with limited energy efficiency criteria in Version 1.0, then building additional detail/complexity in future revisions of the specification.
- Software features such as compression and deduplication should be included in the specification (as minimum requirements or via reporting on the Power & Performance Datasheet), but quantifying benefits will be complicated.

Energy Efficiency Criteria

- Idle state power consumption is a valid proxy for active state, with substantially reduced test burden.
- Idle state should not be considered by ENERGY STAR, since idle is a rare occurrence for Online storage and the nature of disk access in Idle is different than the nature of access in Active.
- Active metrics should be normalized to a “typical configuration” that is representative of an actual customer deployment, to fairly amortize overhead power consumption across all drives.

Energy Efficiency Criteria

- Single-output and multi-output (or, redundant-capable and non-redundant-capable) PSUs require different efficiency targets.
- Recommend additional 0% load test point for informational purposes (no requirements) to better educate purchasers.
- Recommend continued data collection for storage PSUs (especially single-output), to better characterize typical performance vs. Climate Savers / 80PLUS levels.
  - Recommend CSCI Bronze for single-output and Basic for multi-output.
  - Recommend CSCI Platinum for single-output and Gold for multi-output.
- Storage PSUs are often oversized to accommodate power demands of future drive upgrades.

Scope

- The proposed taxonomy does not include an upper limit on some aspects of product complexity.
- Version 1.0 scope should be limited to Online 2 & 3, Removable 2 & 3 categories. Experience & validation of the approach with smaller systems will facilitate expansion to larger systems.
- Version 1.0 scope should include Near-online 2 & 3 (in addition to above).

Taxonomy

- Avoid references to specific technologies (SSD vs. HDD, etc.) to maintain an agnostic approach.
- Taxonomy definitions should be periodically re-assessed to ensure their relevance to ENERGY STAR objectives.
- EPA should consider database engines for inclusion in ENERGY STAR, either as storage, server, or an appliance.
- A concise product family definition is necessary to bound the number of variables that can be assessed, and to minimize testing burden on manufacturers.

Taxonomy

- Taxonomy terms are open to interpretation (e.g. “a portion”, “long term”)
- “Hybrid” should refer to systems that cross taxonomy categories, not systems that use multiple types of disks.
- MaxTTD is defined at the component level, and not the system level, which makes classification of hybrid systems difficult when a variety of devices are in use.
- Definitions for “storage device” and “storage element” are ambiguous.

Proposed Product / Family Approach

Goals

- Identify products and configurations that provide superior energy efficiency
- Fairly and consistently represent energy efficiency benefits of valid product configurations to:
  - end users
  - sales/fulfillment channels
- Minimize testing/reporting burden for ENERGY STAR partners

Assumptions

- For Online storage, disk and controller configuration appear to be the most meaningful factors for energy consumption.
- Software is important, but ENERGY STAR implementation is still under evaluation.
  - At a minimum, software details will be included on the Power & Performance Data Sheet
- Impact of RAS features and Group 3 / 4 differentiation were not obvious from the initial data collection
  - Vendor map of current product lines vs. the proposed taxonomy
  - Additional focused data collection?

Controllers vs. Capacity

- 2 samples submitted
  - F5..F6 → 1 & 2 Drawers
  - I1..I3 → 2, 4 & 8 Drawers
- Inconclusive results
  - Random Read shows large variation
  - F series best performance w/1 drawer
  - I series best performance w/8 drawer

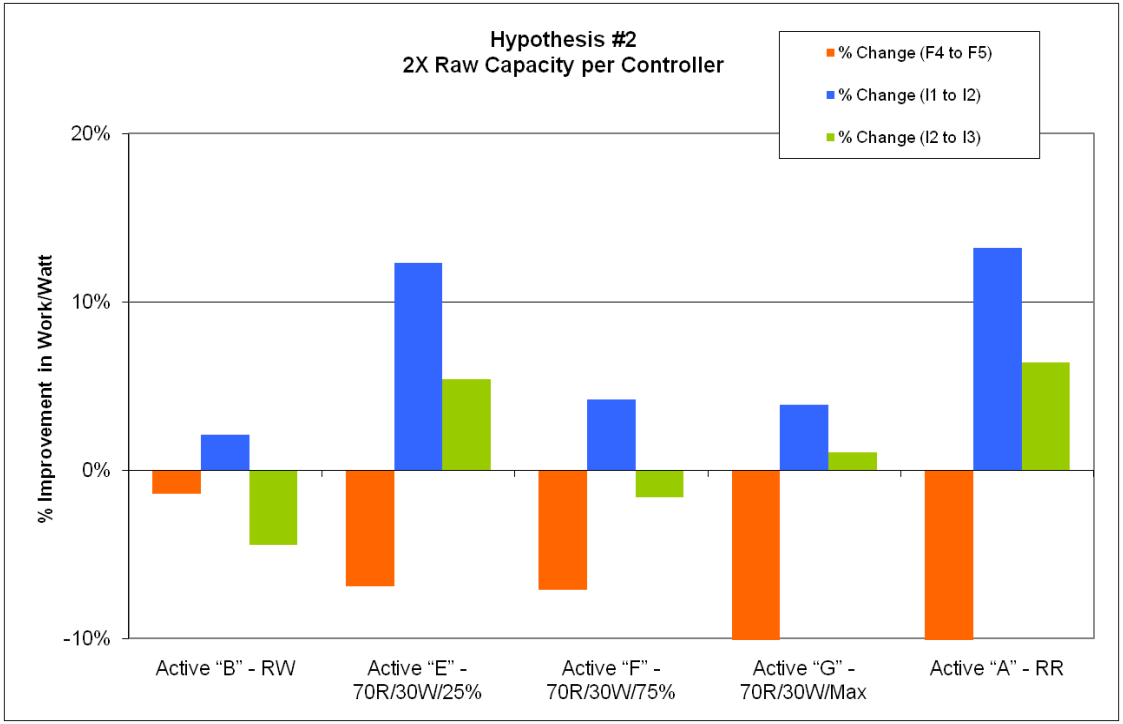

Single vs. Redundant PSU

- Two libraries submitted
  - Reduced efficiency during Active test (both absolute power draw and work / watt)
  - Mixed results for Ready Idle

	Single PS	Redundant PS
Active A		
Single Drive (O1 vs. O2)	54W / 2.23 MiBS/s/W	65W / 1.89 MiBS/s/W
Single Drive (P5 vs. P6)	81W / 1.94 MiBS/s/W	95W / 1.67 MiBS/s/W
Quad Drive (P8 vs. P9)	145W / 2.19 MiBS/s/W	155W / 2.03 MiBS/s/W
Ready Idle		
O systems	37W	54W
P Systems	23W	23W

Product / Family Proposal

"A fully-functional storage system that supplies data storage services to clients and devices attached directly or through a network… A storage product composed of integrated storage controllers, storage media, embedded network elements, software, and other devices…"

- Common architectures:
  - Online Storage: Controllers and Drawers
  - Tape Libraries: Integrated Drives and Robotics
- Other architectures:
  - Highly-integrated storage – 'full rack' modularity
  - Server-based storage
  - Hybrid storage solutions
  - Storage appliances

EPA

Product / Family Proposal

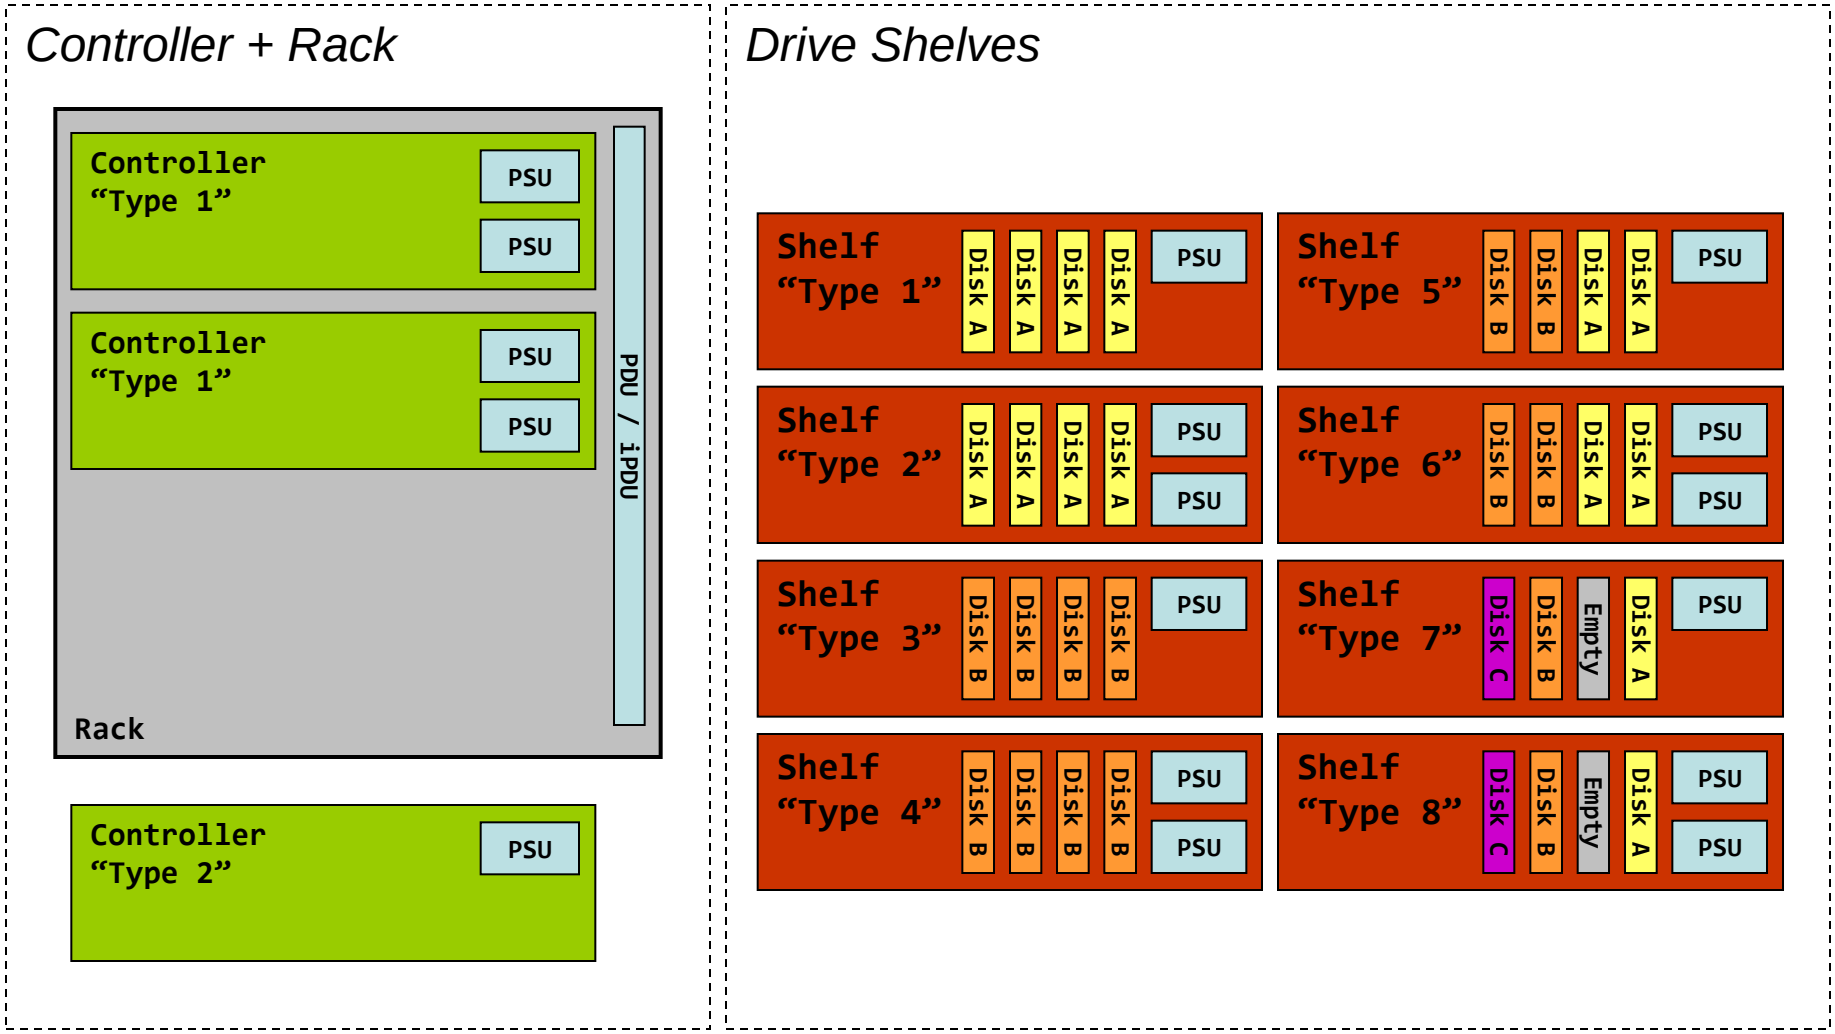

Product / Family Proposal

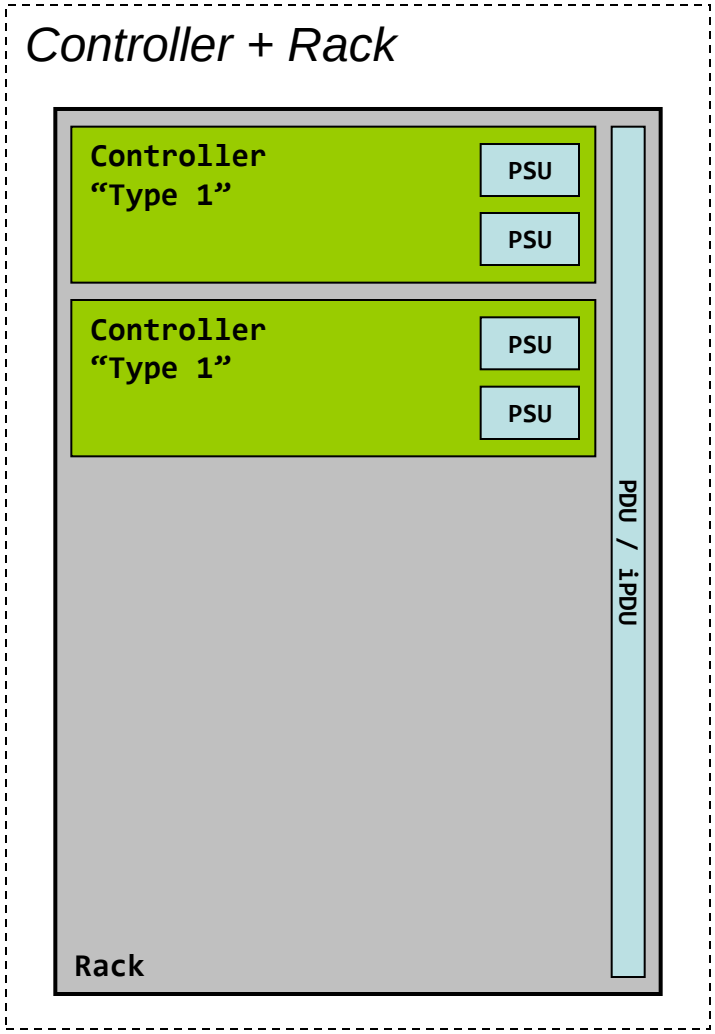

- Rack + Controller configuration defines the “Base Type” for the family
- Any substantial changes are considered a unique family
  - Controller hardware / software configuration
  - Quantity & redundancy of controllers

Product / Family Proposal

- Unique “Shelf Types” are modular elements to be added to the Base Type.
- Drive shelf type is defined by disk type & combination, and PSU load.

Drive Shelves

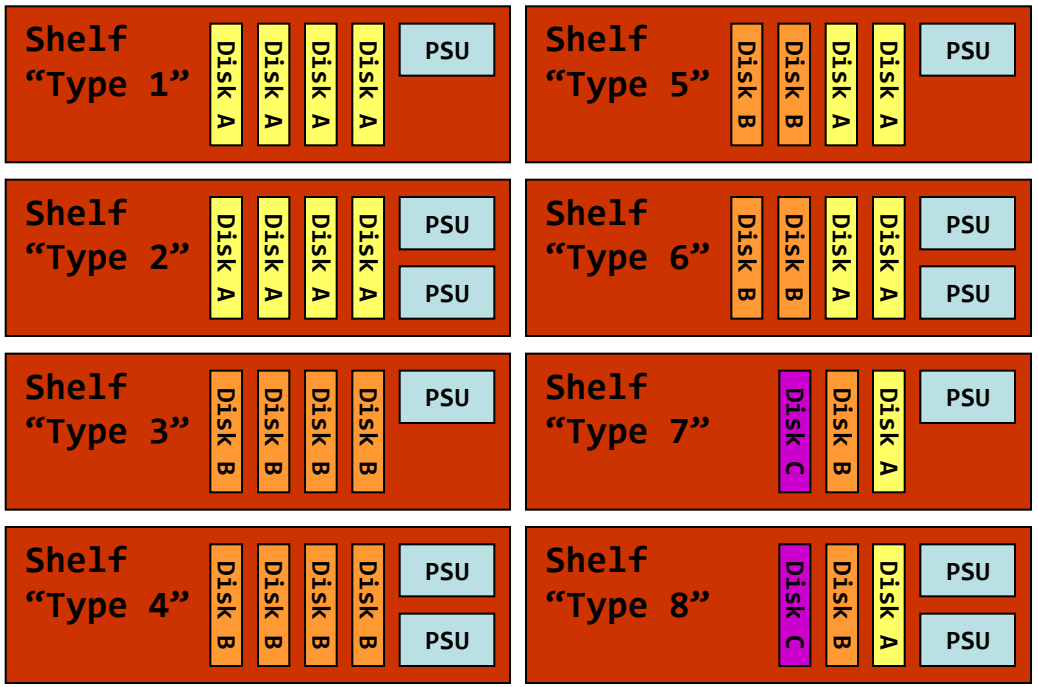

Product / Family Proposal

- Test Base Type with n and 2n “fully-populated” samples of any Shelf Type.
  - If qualified with “n”, and energy performance improves with 2n, then qualification extends to larger systems at the vendor’s discretion.
  - Vendor must submit “technical justification” of max configuration.
  - If energy performance does not improve from “n” to “2n”, test to identify max configuration that does qualify.
- Test same Base Type with a different Shelf Type.
  - If qualified, qualification extends to Base Type + all quantities up to max capacity for each type and all combinations
  - Compute quantity of each type of disk from max system size to establish max quantities allowed in blended configurations.

Qualification Testing

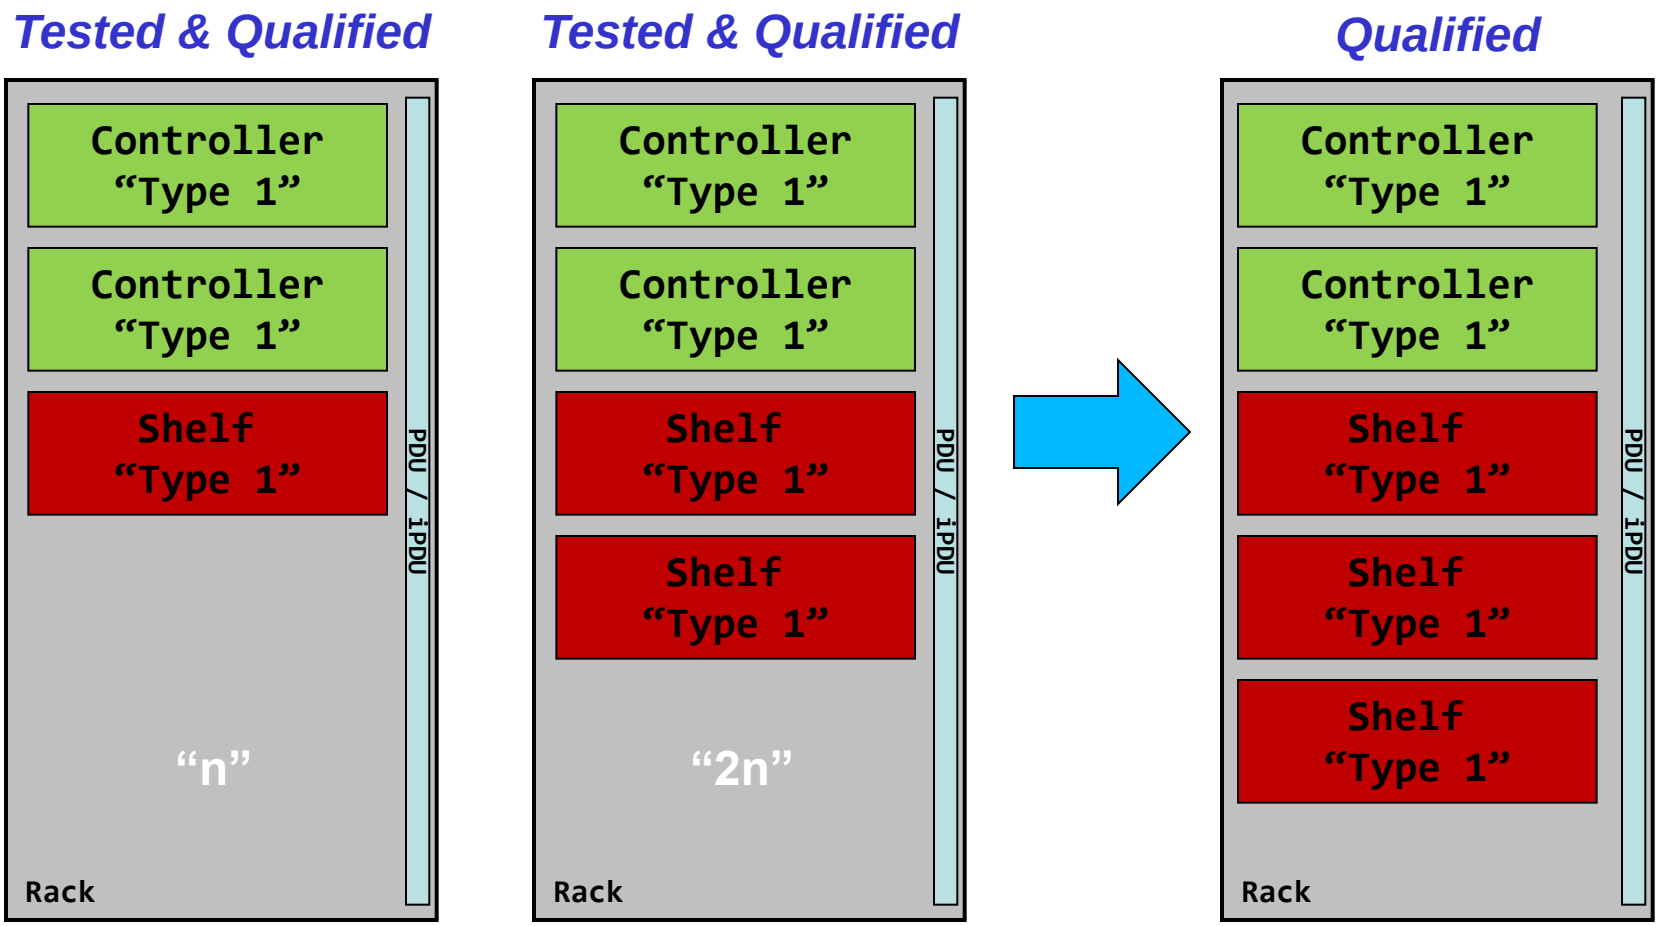

Qualification Testing

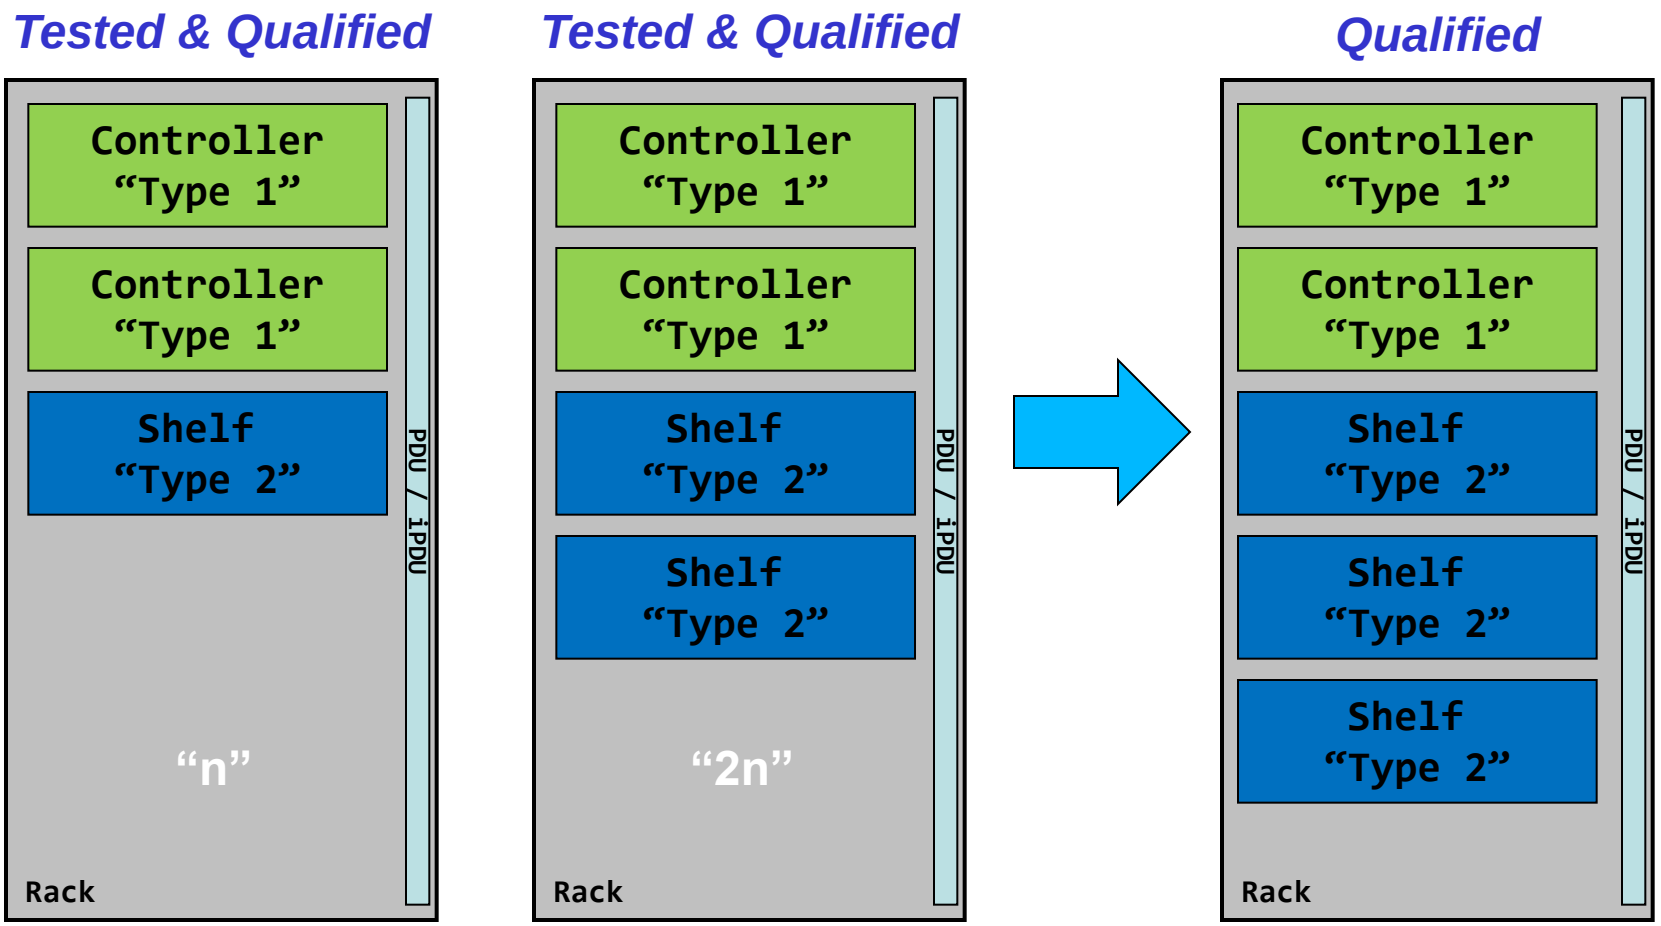

Qualification Testing

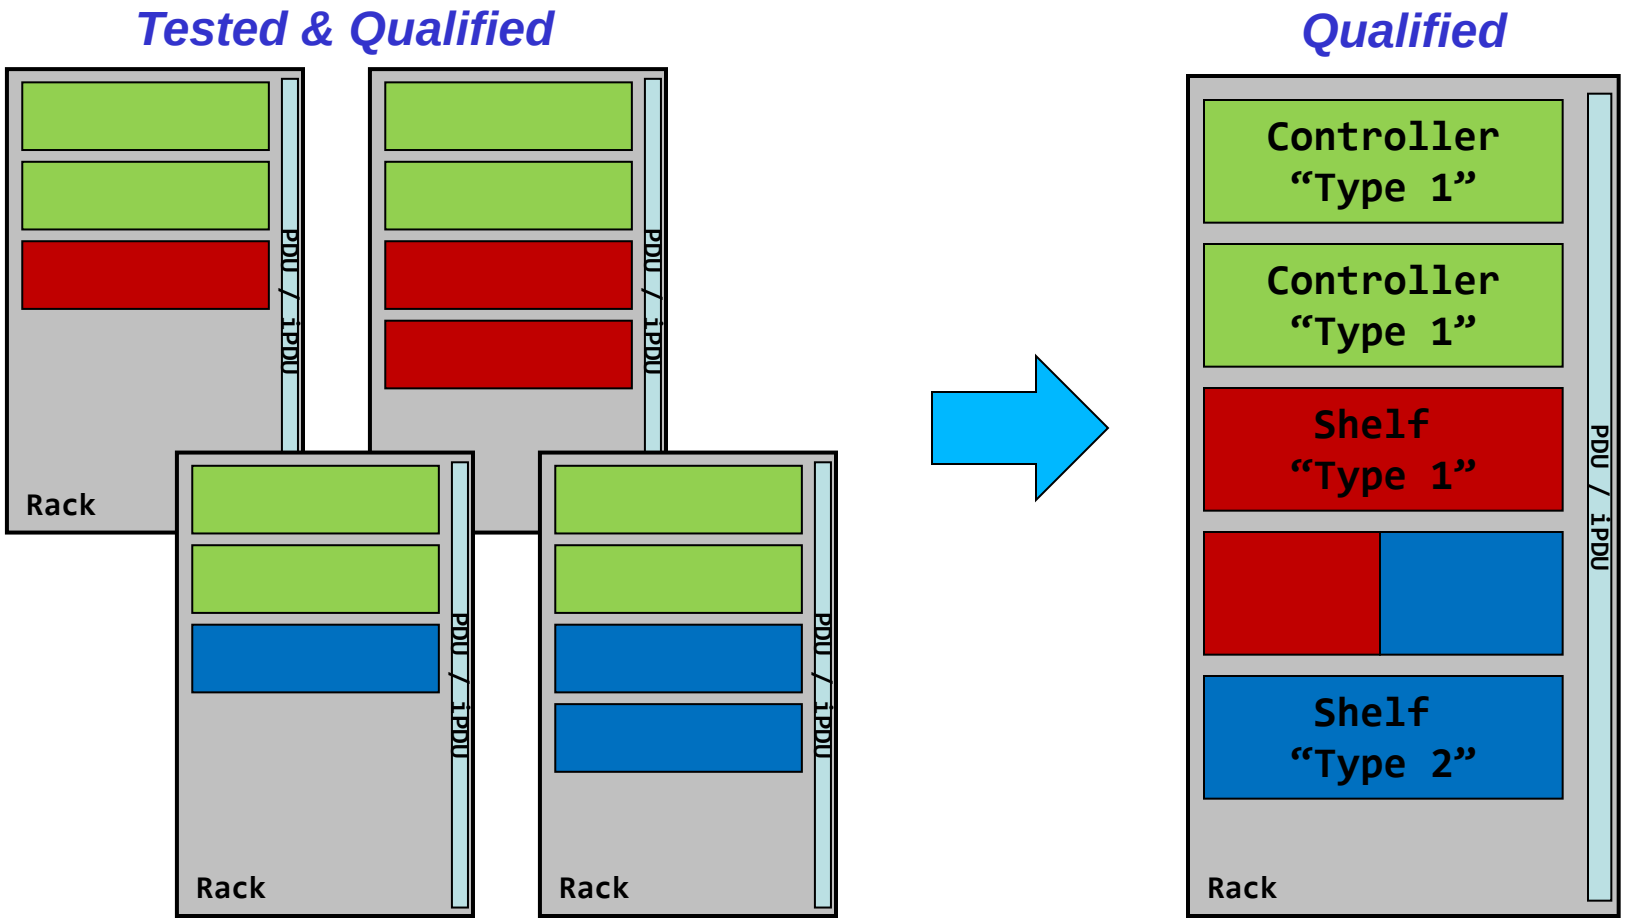

Product / Family Questions

- Should “n” = 1 or be at the vendor’s discretion?
- What constitutes a “fully populated” drawer?
- If Base Type tested and qualified with redundant controllers, does that automatically qualify the single-controller configuration?
- If Shelf Type tested and qualified with redundant PSUs, does that automatically qualify the single-PSU configuration?

Product / Family Questions

- Does this approach to product families reasonably match vendor fulfillment models (e.g. Rack + controller configuration = basic module; drive shelf types = extension to basic module)?
- How does software factor in to the Base Type / module concept?
- Will this qualification approach be meaningful to end users?
- Can a similar approach be designed to apply to tape libraries?
- What “product” earns the label and how will it be represented to purchasers?

Frame vs. Drives (Tape)

- Two libraries submitted. Each showed better work/watt as tape drives are added to the system.

	1 Drive	2 Drive	3 Drive	4 Drive
Active A				
2 Drive Frame (P2 .. P3)	1.6 MiBS/s/W	2.0 MiBS/s/W		
Quad Drive Frame (P4, P5, P7, P8)	1.5 MiBS/s/W	1.9 MiBS/s/W	2.1 MiBS/s/W	2.2 MiBS/s/W
Ready Idle				
2 Drive Frame	37W	Not recorded		
Quad Drive Frame	23W	23W	23W	23W

Round 2 Data Collection

- Focus from preliminary data collection:
  - Impact of RAS features
  - Performance differentiation of Group 3 & 4
- Focus from product / family hypothesis:
  - “n” and “2n” configurations
  - Single- vs. dual-controller
  - Redundant disk shelf PSU vs. non-redundant
  - Mixed drive types in drawers
  - Mixed drawer types in rack

Development Schedule

- *6/4/09* *Framework distributed*
- *7/20/09* *Stakeholder meeting (San Jose)*
- *10/15/09* *Test Procedure Workshop (Phoenix)*
- *12/28/09* *Start 1st round data collection*
- *2/2/10* *Stakeholder meeting (San Jose)*
- *3/1/10* *Complete 1st round data collection*
- *4/9/10* *Draft 1 distribution*
- *4/15/10* *Stakeholder meeting (Orlando)*
- *5/21/10* *Draft 1 comments submitted to EPA*
- *7/20/10* *Stakeholder meeting (San Jose)*
- *Jul-Aug* *Supplemental data collection*

Contact Information

- *Una Song (US EPA)*
  - song.una@epa.gov // 202.343.9024
- *Steve Pantano (ICF)*
  - spantano@icfi.com // 202.862.1551
- *Al Thomason (TBWC)*
  - thomasonw@gmail.com

More Info:
http://www.energystar.gov/NewSpecs