



Education

# **How to Choose and Evaluate Storage for your Environment From the Consumer Point-of-View**

Chris Lionetti, Microsoft Windows Live Operations

- The material contained in this tutorial is copyrighted by the SNIA.
  - Member companies and individuals may use this material in presentations and literature under the following conditions:
    - ◆ Any slide or slides used must be reproduced without modification
    - ◆ The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
  - This presentation is a project of the SNIA Education Committee.
  - Neither the Author nor the Presenter is an attorney and nothing in this presentation is intended to be nor should be construed as legal advice or opinion. If you need legal advice or legal opinion please contact an attorney.
  - The information presented herein represents the Author's personal opinion and current understanding of the issues involved. The Author, the Presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.
- NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.**

- **How to Choose and Evaluate Storage for your Environment From the Consumer Point-of-View**
  - ◆ This presentation will describe a repeatable process for evaluating storage product to ensure that your business needs are met, and your operational risks are minimized.
  - ◆ This selection and evaluation process will allow you to select best of breed and go from a poorly defined request for storage to a vetted/tested/defined/deterministic successful deployment.
  - ◆ At the end of this presentation you should have a well defined process to follow with clear milestones, a clear list of defined interactions with vendors, and way to change your vendor relationship from a buyer/seller to a vendor/partner.

- This presentation focuses on II repeatable steps
  - ◆ Know Thyself
  - ◆ White Paper Analysis
  - ◆ Requirements
  - ◆ Initial Scoring
  - ◆ Test Plan
  - ◆ Acquire equipment
  - ◆ Test Phase
  - ◆ Document Results of Testing
  - ◆ Determine Proper Scaling Unit and CAPEX/OPEX
  - ◆ Make Recommendation with Confidentiality/ Transparency
  - ◆ Pilot installation and feedback to stage I

# Step I : Know Thyself

- You need to have a grasp of what you currently do business on. What is the IOPS of a **specific** application, what is acceptable? What is the latency of that app, how much of that is due to Disk versus CPU versus Memory
  - ◆ Document the weaknesses of your current infrastructure.
  - ◆ Document its strengths.
  - ◆ Document the limiting factors that prevent you from scaling what you have to handle more load.

- Steps taken on this step set the stage for all following steps.
  - ◆ Spend time interviewing the application owners.
  - ◆ Ask pointed questions. Get direct answers
    - › Do they need more storage OR faster storage. Make them prioritize between the two.
  - ◆ Capture performance from existing equipment to determine how well they utilize what they have.
  - ◆ Spend time observing storage consumer workflows.
    - › Draw Sticks and Bubble diagrams of your environment
  - ◆ Figure out where data flows and where networked storage can give you biggest bang-for-you-buck.

- Document Baselines for all values you expect to judge new storage on. These are control numbers for testing
  - ◆ What are your current Gb/Watt, Gb/RackU, Gb/\$
  - ◆ What are your current IOPS/Watt, IOPS/RackU, and IOPS/\$.
  - ◆ What are your current (MB/s)/Watt and (MB/s)/\$.
  
- Document your CapEx and OpEx for current infrastructure
  - ◆ OpEx can be excessive as equipment ages. Power, Warranty uplift charges, and Job Task Analysis should be considered.
  - ◆ CapEx should be considered if your need to expand what you already deploy. You should know better than anyone what your growth rate is.

# Step 2 : White Paper Analysis

## ➤ The grass is always greener....

- ◆ Consider your current infrastructure Weaknesses
- ◆ Consider your current scale and storage growth
  - Are you about to hit a wall with your existing infrastructure? This is common with DAS to SAN conversions.
- ◆ Consider how you could optimize the workflow to remove bottlenecks
  - This could be as simple as offloading backups to a dedicated server using snapshot technology and preventing your backup traffic from traversing an Ethernet structure.
- ◆ Consider new compliance rules that you are being asked to support
  - Will your current solution support encrypted LUNs? Do you expect your General Purpose Processors within your servers to do that work, or would you offload than to encryption devices.

# Step 2 : White Paper Analysis

## ➤ Don't Eat the Yellow Snow....

- ◆ Remain Diligent and skeptical to radical changes to your environment
  - This can be a one way street on a crazy car driven by your vendor
- ◆ Avoid sales people that present solution to your problems without knowing what your problem is.
  - When the only tool you have is a hammer, every problem looks like a nail.
- ◆ Avoid the vendor who presents a competitive TCO that has no pricing details and makes huge assumptions about your infrastructure.
  - I have had a vendor state they could save me 50% on my storage. Didn't ask me what I currently pay, nor what vendor I currently use, or the type of storage I deploy. After 3 weeks of trying to wrangle out of the vendor a rough estimate of price, it turns out he was more than 2x the price of my current storage.
- ◆ Beware the 'That is on the Roadmap' list.
  - I have had a vendor offer to sell me today's storage that doesn't meet my needs and then swap out when the new product releases. This would lock me into the new storage without testing, and would force me into storage today that doesn't meet requirements.

# Step 2 : White Paper Analysis

## ➤ Message to Storage Vendors ;

- ◆ STOP SHOWING ME THE &\$@#ING MAGIC QUADRANTS.

- If you promise that it will be a Magic-Quadrant-Free meeting, I will promise that I won't wear your competitors T-Shirt when I meet your CEO/CTO/Engineering staff.

# Step 3 : Requirements

## ➤ Define your stakeholders.

- Stakeholders are the people that will be helped or hurt by the new solution. These could represent your day-to-day operations team, your business or application owners, or your procurement and security teams.

## ➤ Create a Skeletal structure for your requirements

- Insert sections for each subgroup of Requirements
- Insert Requirements derived from knowledge you gained in Step 1 of this process

## ➤ Stakeholders create and modify the requirements.

- Q: How did NASA get to the moon on a rocket with millions of parts all made by the low bidder?

# Step 3 : Requirements

- Sample Requirements subsections might look like this, but a full requirements document will be hundreds of items.

Operation	
	Array must have an event log that can be either dumped or exported to a CSV or XML friendly format.
	Array must have a Priority in the event log that can be configurable to escalate critical events above informational noise
	Array must report common SCSI errors such as SMART, SoftSCSI, and Heroic-Recovery errors
	Array must report common loop faults
	Array must have Power LEDs to identify the functional status of all FRU components
	Array must have fault LEDs to identify the fault status of all FRU components
	Array must have Hot Swap Capacity to replace all Active Components. Active is defined as contains mechanical or IC components. A Backplane that has no ICs is not considered active.
	Array must support a method to register multiple HBAs from a single server and manage them as a group
	Array must have N+ 1 capability for all active components.
	Array must have hot code load
	Array must have hot application load
	Array must support an online conversion from Direct Attached to SAN Attached.
	Array should be able to support an outage of 72 hours without data cache loss. De-stage is preferred.
	Array should not require intervention during a power restore.
	LUNs should support a LUN Expansion process.
Reporting	
	Array must allow for a template application of Phone Home/Email Home settings to be applied
	Array must support SNMP trap reporting 1.1 and 1.2 and forwarding
	Array must support SMI-S 1.2 in a read mode.
	Array must support a connectivity screen that shows which servers are being serviced, and which are active
	All critical events presented by the array should be English readable and actionable.
	All events should have unique ID tags and Timestamps
	All events should be stored on the array while also allowing forwarding to an event monitoring server
	Must have an option to purchase an SRM reporting
Security	
	Array must integrate with a LDAP/AD/Radius account management system
	Array commands via CLI and GUI must be encrypted (e.g. SSL or SSH).
	All commands to the Array should be logged including the issue credentials of the requester
	No In band management of array is allowed, and no storage access should be allowed via out of band port.
	No default unchangeable passwords should exist to manage the array (unchangeable service accounts)

# Step 3 : Requirements

## ➤ Define your Priorities.

- ◆ Pri 1 Requirements : These are Requirements that you will NOT adopt without. If the array doesn't support this feature, then you see no point adopting this product.
  - Be careful using this priority as it acts as the ultimate veto.
- ◆ Pri 2 Requirements : These are features that you feel should exist on the product. If you encounter a threshold of enough Pri 2s that are unmet, you will pass on a product and not adopt it.
- ◆ Pri 3 Requirements : These are features that you would like to see, but they are not alone or in combination capable of removing a product from consideration.
  - These features are items that will be used to separate the products being evaluated via scoring.

## ➤ Define your Weighting

- ◆ A weight is a relative value of a feature. All Pri 2's and 3's should be assigned a weight from 1 to 10.
  - Two features with weights of 4 are equal to a single feature that has a weight of 8

# Step 4 : Initial Scoring

- Using your Stakeholder Approved Requirements which has been Prioritized and Weighted. You can proceed to research vendors and products.
  - ◆ Start by evaluating each product against your Pri 1 Requirements. This will eliminate most vendors and leave you with a list below 10 who may be able to compete.
  - ◆ Of the 10 or so vendors left. Evaluate the Pri 2's and Pri 3's. Leave blank any requirements that you cant derive from the product datasheets.
    - You may find that the more data you fill in the closer you can get to a final score.
    - As an example, I might find that I have 735 points possible if all my weighted requirements are met.
      - » Vendor A scores 395, with a possible 250 unknown.
      - » Vendor B scores 510, with a possible 150 unknown.
      - » Vendor C scores 210, with a possible 175 unknown.

## ➤ Narrow the Field

- Now that you have a list of direct pointed questions about the products that a vendor makes. Approach that vendor and fill in those unknowns.
- ◆ At this point, You may be left with questions on the requirements that deal exclusively with pricing versus IOPS.
  - It would be prudent at this time to obtain ‘Not To Exceed’ type pricing to ensure you are dealing with the correct class of array.
- ◆ At the end of this process you should have;
  - Handful of vendors, that meet you Pri I requirements.
  - A Score for each product that is specific to YOUR business needs.

## ➤ Who recognizes the source of the

- ◆ **Five Stages of Acquisition:**
  - > Infatuation
  - > Justification
  - > Appropriation
  - > Obsession
  - > Resale

# Step 5 : Test Plan

- **caveat emptor – It is your responsibility to ensure that your requirements are met.**
- **You alone know how you will use a feature.**
  - › You may have a requirement that Events concerning replacement parts all be human readable and actionable. Technically Hexidecimal is Human readable assuming that human has the correct mystery decoder ring. While a vendor may mark this as pass, I would mark this as a fail.
  - › You may have a requirement to stock spare parts on site. You may have assumed that the vendor understood that you might want to replace a failed component yourself. If you have to wait 8 hours for a technician to type a specific code to allow that spare part to be seen as fixed it may defeat the purpose of the local spare parts.
    - The devil is always in the details, and those individual details may make the implementation of a specific feature that you need/want impossible

## ➤ Develop a IEEE-829 Standard Test plan.

- Your test plan should outlay how you plan to set up your lab. What equipment you will have available. How you will cable up the individual hardware to the unit being tested. Versions of all hardware (HBA's Switches, Servers, OS, Arrays) including patch levels and firmware.
- Your test plan should also outlay the load generation tools you plan to use, what profiles you plan to employ, and how you plan to monitor and evaluate the performance. This should mimic the information you discovered in step 1 as closely as possible
- Each Requirement should be track-able to a test case or set of test cases. These cases are designed to prove the features of the device against the requirements document.

- **Develop a IEEE-829 Standard Test plan.**
  - ◆ **Each Test Case should have the following sections.**
    - Setup Pre-Conditions
    - Test Procedure
    - Pass/Fail Criteria
    - Expected Results Versus Actual Result
    - Log for Time / Date / Runner
  
  - ◆ **Define at least 4 levels of priority – Don't cry wolf**
    - Severity 1 - The discovered issue causes data corruption
    - Severity 2 – The discovered issue could cause corruption, but can be worked around, or avoided, or limits functionality but does not endanger data
    - Severity 3 – Issue does not corrupt data or limit functionality in any way, but it just isn't right. Such as a button on the GUI doesn't work, but the CLI does.
    - Severity CR – Change Request. Technically not a bug, but also non-intuitive or could be done much better.

# Step 6 : Acquire Equipment

- A common ask from a vendor when requesting equipment for test is for you to sign an evaluation agreement, and a Zero Dollar PO. You will likely also be asked to sign an NDA.
- ◆ Have your legal department review this documentation before signing anything.
  - I have had a vendor slip a \$50K+ installation fee into an evaluation such that if we adopted the hardware it would be paid by the vendors sales organization so we would never be billed, but if we decided to not adopt the storage it would be directly charged to us. Good news – Legal caught it.

- Test, test, test, test....Oh my god, make it stop...
  - ◆ Run your test plan through first to ensure that all Priority 1 tests succeed. Then continue on to Priority 2, then finally Priority 3.
  - ◆ I always phase my tests so that if I encounter a Higher Priority unmet requirement I have longer to address it and accept a fix.
    - It is very common to fail a test due to a misunderstanding of how a feature works. Don't feel bashful asking why feature X doesn't work and being corrected on how it is supposed to work. Every vendors storage is like learning a new language.
  - ◆ All tests that fail the first time through that have either been patched, or you now understand how to operate that feature can be re-run or regressed.

## ➤ How often to Escalate Issues.

- ◆ Consider breaking the test cycle up with a vendor and escalate in the following manner.
  - › With a Priority 1 Requirement that is supposed to pass but doesn't, or a Severity 1 Test Fail gather all logs and escalate immediately to the vendor.
    - Since the test plan already outlines all versions and a diagram of the environment, include that with the logs and contact the vendor representative/engineering contact.
  - › If the failure is a Severity 2 failure as defined in the Test Plan, still gather all the logs, but will wait until the next scheduled weekly meeting and bring these up with the other Severity 2 issues.
  - › If the issue is a Severity 3 failure, you can usually document these after the first complete run of the test plan, right before you start my regression run.
  - › A Severity CR Change Request is usually only documented on my final evaluating document as it will not impact the requirements score.

# Step 8 : Documenting Results

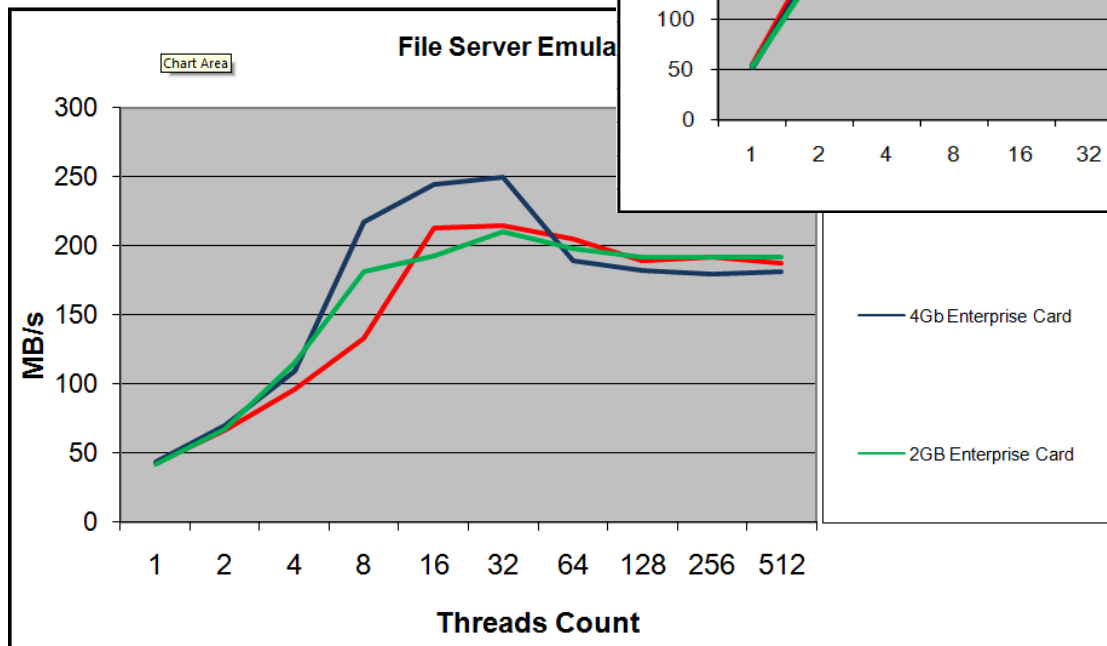
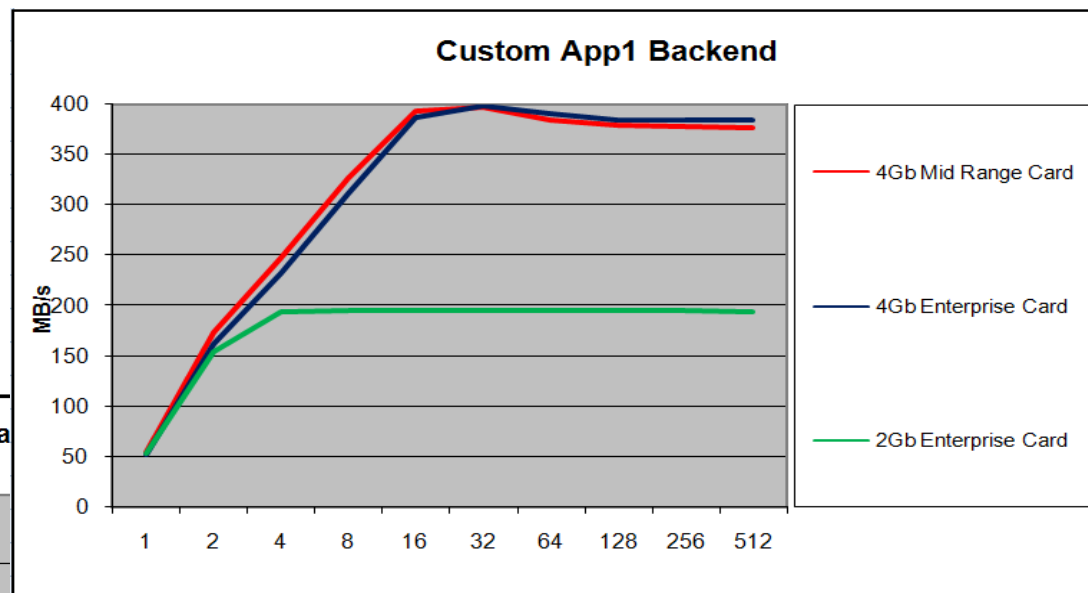
- This is going to wrap all previous steps into a concise document.
  - This document will discuss your current environment, usage models and job task analysis of how you currently spend your Operational time, a definition of what problems you are trying to solve.
  - This document will also discuss your requirements, how each product plans to solve those requirements as well as your testing results that expose the disconnects between requirements and real life usage.
  - Additionally you are going to want to document procedures such as the installation and common configuration steps. This document can either point to or become a deployment guide.

# Step 8 : Documenting Results

- Many of your results be better represented in graph form or picture form.
  - ◆ i.e. To determine which HBA to use when upgrading from 2Gb/s cards to 4Gb/s cards. The choice was to go with an enterprise card like the 2Gb/s card we used already, or go with a mid range card.
  - ◆ The requirement reflected a need for a new card to replace an old card but also reflected the need to keep costs constrained.
    - One of the Priority 2 Requirements was that the new 4Gb cards outperform the 2Gb Enterprise card on all tests
    - A Priority 3 requirement was that the cost for the Enterprise 4Gb card be completely justified by increased performance. Ie If the 4Gb Enterprise card is 78% faster, but only 55% faster it would fail. If the card is on average 25% faster but only 10% more expensive it would pass.

# Step 8 : Documenting Results

➤ Sample Graphs of these cards against each other.



# Step 9 : Determine Scale Unit

- Your next step is to determine the sweet spot in terms of the scale unit.
  - The sweet spot will be determined not only by the architecture of the device being tested, but also by the way your application interacts with the array.
- ◆ Some loads are notoriously hard to solve, while others are very simple.



**Check out SNIA Tutorial:  
Solid State Storage for the  
Enterprise**

- You will want to create a model of a number of different scale units with the end result being user transactions per Dollar (or Watt) that fall within your predefined acceptable performance metric.

## ➤ Confidentiality

- ◆ Whether or not you sign an NDA, you should consider your results and evaluation as confidential and possibly competitive information
- ◆ You will create a Master Evaluation may contain proprietary company information, or may contain results data from multiple competing vendors.
- ◆ This Master Evaluation will be what is presented to the Stakeholders and either outline or guide your future storage decision.

## ➤ Transparency

- ◆ Expectations - Vendor might get access Confidential master version.
- ◆ Vendor should be informed about HOW he is and isn't allowed to use your testing results.
  - Is vendor planning on using you as a reference customer? Will you see your name in press releases about the vendors products? How much exposure do you want as a poster child?
- ◆ Vendor should be presented with a cleansed vendor specific version of the Evaluation/Recommendation.
  - Vendor specific version should have all competitive information striped from it.
  - If you tested 3 vendors against each other, you would need to produce 4 documents. A Master version, and 3 Vendor specific versions.

## ➤ Rookie Mistake Alert

- ◆ If you create your vendor cleansed version of your document from your master confidential version, be sure to remove all hidden data.
  - » <http://office.microsoft.com/en-us/help/HA100375931033.aspx>
- ◆ **Encrypt the content of the Document. Consider setting the ‘no content extraction’, ‘no printing’.**
  - It wont stop a determined person, but it will make your intentions known.
- ◆ **Digitally sign your document.**
  - That digital signiture will become invalid should the content change. Also ensures that a vendor isn't misrepresenting your data since you can know if it came from you.
- ◆ **Consider exposing the vendor to there document before submitting the master document.**
  - Request rebuttal type feedback

## ➤ Balance

- ◆ A Balance between Transparency and Confidentiality is possible.
- ◆ The need for Transparency may be outlined as a requirement to get loaner equipment.
- ◆ Transparency is a Win-Win for the Industry and the consumer.
  - When the vendor roadmaps match the true needs of the datacenter we all get products we need instead of what they try and shovel down our throats.
  - Green Datacenter is a perfect example. The is only a selling point when I as the customer demand it. It doesn't directly put a thin dime into the vendors pocket, but it does show up on my requirements as a OpEx threshold.
  - TELL THE VENDOR WHY THEY DON'T MEET REQUIREMENTS. Together we have a voice.

# Step 11 : Pilot Installation

- Make sure you have a playbook before a single screwdriver is picked up.
  - › Know what you plan to do before you do it.
  - › Review the playbook to ensure that it fits in your process.
- Understand the Warranty/Uplift/Site Services responsibilities.
- Dry run the support process.
  - ◆ Burn in the array, and simulate a failure.
    - › Test Site Escalation response times. Test Vendor response times

## ➤ Document Everything

- ◆ Keep an eye out for anything out of the ordinary.
- ◆ Review at minimum at 30 day, 60 day, and 90 day.
- ◆ Interview Ops people about challenges and benefits.
- ◆ Update both the ‘Know thyself’ step as well as create new items for the ‘Requirements’ step.
  - › This is the iterative process steps.

- Please send any questions or comments on this presentation to SNIA: [trackstorage@snia.org](mailto:trackstorage@snia.org)

**Many thanks to the following individuals  
for their contributions to this tutorial.**

**- SNIA Education Committee**

**Robert Peglar**