Economizing on Storage Infrastructure

Why Storage Costs So #!@$ Much and What You Can Do to Bend the Cost Curve

Presented By

Jon Toigo
jtoigo@toigopartners.com

Toigo Partners International

IT-$ENSE

DMI Data Management Institute

How's that Storage Workin' Out For Yah?

Storage now accounts for between .33 and .75 of every $1 spent on IT hardware

And That's Just Acquisition Cost

For every one of these that you spend on a disk array…

…your actual annualized cost of ownership is ***5-8x*** that amount.

Backup
Maintenance
Administration
Environmentals

Acquisition

No Problem, If You Have

A set of these…

Or One of These

Pretty soon, you are talking about real money…

So, Why Does Storage Cost So Much?

- Disk drives aren't that expensive…
  - Everyone is just selling a box of drives from the same folks, aren't they?
  - And drives continue to double in capacity ~ every 18 months, while dropping 50% annually in $-per-GB…
  - And OEMs mostly use the same shells (but with different faceplates and logos – extra for black)…
  - And RAID dates back to 1977…

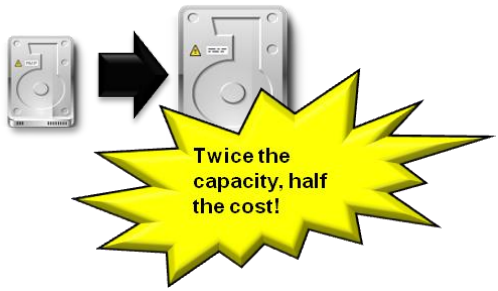

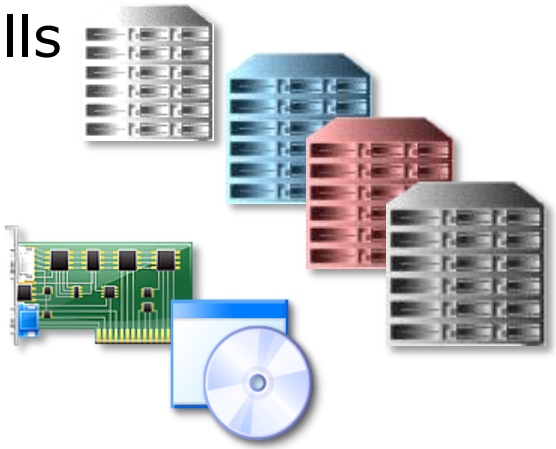

Baseline Cost Model

COST TO BUILD

- Raw Materials/Components
- Assembly Costs
- Value-Add Software
- Supply Chain Management
- Direct/Indirect Channels
- Distribution Costs
- Warranty and Maintenance
- RMI Churn
- Legal Fees and IP Protection
- Hedging

What is "Value-Add" Software?

- Differentiation of otherwise commodity hardware requires **value-add** functionality
  - Publicly-traded OEMs must "innovate" (aka "add value via software embedded on array controller") **every 6-8 months**, or risk losing market analyst interest
  - Embedded software used to **lock in customer/lock out competitors**
  - Portrayed as "one-stop-shop" or "pre-integrated" or "intelligent or smart(er) product" with **TCO reduction value**
  - **Software costs nearly $0 to mass produce** – a plus

Joining more proprietary software to an array controller encourages "forklift upgrading" by customers

For Example

- A popular de-duplication solution vendor offers two kinds of de-duplication rigs
  - A. Software only
    - Licenses software separately for $2500 per TB
    - Run on a server, use any storage you have
  - B. Software plus hardware
    - 32 TB of commodity SATA disk storage with embedded software = $410,000 MSRP (or $12,812.50 per TB)
    - 1 TB SATA disk costs $79 - $140 on NewEgg.com (~$3200 for 32 TB)
    - Vendor's "value-add" rig costs $80K for software, $330K for disks or **100x the cost of the disks themselves**

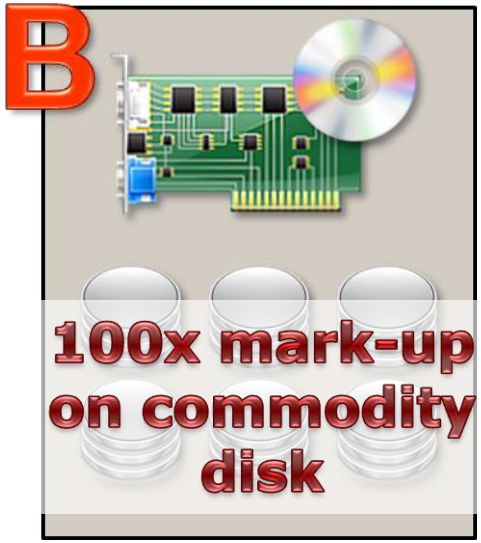

But, the Vendors Claim…

- You get 70% reduction ratios with their rig, so…

"***Phenomenal cosmic powers.***
Itty, bitty living space."
Genie from *Aladdin*

(We like fairy tales, too.)

Another Case

- Thin provisioning: "virtualized" capacity oversubscription, but with "intelligent" demand forecasting
- Vendor claims better capacity allocation efficiency
- Starter rig:
  - 8.7TB (raw) mix of (6) 460GB 15k SAS and (6) 1TB 7200 SATA drives at $38.1K with single controller and value-add software for tiering and thin provisioning.
  - $2600 worth of drives for $38.1K

6 SAS Drives ~$350 each without OEM volume purchase discount ($2100)

6 SATA Drives ~$85 each without OEM volume purchase discount ($480)

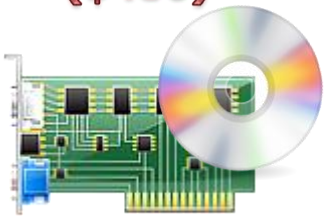

Head, chassis, software, margin = ~$35.5K???

But, Doesn't Thin Provisioning Help Capacity Management and Save Me from Buying More Disk?

- Sort of…

But to make this shell game reliable…

Probably best to keep a stack of these handy as a buffer against "Margin Calls"…

(Guess you still need that spare capacity…)

Think About This One

Brand Name NAS

- 6TB (4TB usable) Seagate drives in RAID 6
- Proprietary OS
- CIFS/NFS/iSCSI/FCP licenses
- 3 year NBD support
- "Heavily discounted" at $1.56 per GB raw or $9,604

Off-Brand NAS

- 24TB (16TB+ usable) same drives in RAID 6
- Linux OS
- CIFS/NFS/iSCSI/FCP licenses
- 3 year NBD support
- No discount $.70 GB raw or $17.2K for 4x usable capacity (other vendor discounted after learning of this bid)

"Staple of enterprise-class file storage"

"Staple remover"

Necessary Caveat

- Most storage isn't purchased like ground beef (by the pound), of course.

- That said, why pay *filet mignon* prices for *flank steak*? (Especially when user files tend never to be accessed again within 30 days of being written to NAS disk.)

Clearly, Some Functions Need to be Performed at the Array Level

- Configuration controls, box monitoring and management…*um*…*eh*…*phone home?*
- But why these features?
  - Intra-array volume replication and snapshot
  - Array to Array Replication
  - On-box Automated Tiering
  - Thin provisioning
  - Deduplication
  - Point in Time Split Mirroring
  - Content Indexing/Content Addressing
  - (insert latest feature here)

Other Ways to Build Storage

- Why not place "value-add" functions where they can be shared instead of stove-piping them?
  - Storage virtualization
  - On gateways or routers
  - On servers
- Less expensive, more extensible, more manageable, and capable of granular provisioning to specific data

Goes to Your Definition of Value

- Vendors claim that "value-add" software
  - Reduces time required for integration (so you have more time to play *World of Warcraft*)
  - Simplifies management by optimizing resources automatically (so you can do more with fewer staff)
  - Demonstrates your "eco-consciousness" (so you can say that you are doing your part to stop global warming)
  - Increases your "cool" quotient (ownership of the latest gear gives you "geek bragging rights" that in turn make you more attractive on date night)

Actual Value (and Cost) are Partly a Function of Channels

COST TO BUILD

- Raw Materials/Components
- Assembly Costs
- Value-Add Software
- Supply Chain Management
- Direct/Indirect Channels
- Distribution Costs
- Warranty and Maintenance
- RMI Churn
- Legal Fees and IP Protection
- Hedging

Channel Basics

- OEMs sell through channels
  - "Direct" – via sales staff
  - "Indirect" – via a complex mix of
    - Vendors (re-branding OEM gear)
    - Value-Added Distributors/Resellers (VADs and VARs)
    - Resellers/Solution Integrators
- Everybody takes a cut

Cost to build array	=	x
Profit margin added by OEM	=	x + (x * ~.25)
Profit margin added by Vendor	=	OEMx + (x*~3)
Profit margin added by VAR	=	Vendorx + (x*~.20)
Profit margin added by Reseller	=	VARx + (x*~.25)

A 20 TB array costing $5,200 to manufacture carries a minimum MSRP of $29,125

Of the Lot, I Prefer a Competent Solutions Integrator

- A good solutions integrator is worth his/her weight in gold
  - More in tune with your business requirements and objectives
  - Supplements your staff, supports decision making with real product experience (knows where the bodies are buried)
  - Can test configurations for you before you buy anything
  - Can assist in integration
  - Provides first (and sometimes second) tier support
- So where are the Jedi Masters?

"No Job Have You"

- In February 2001, the bottom dropped out of storage -- previously viewed as "unflappable"
- Major OEMs and Vendors savaged their indirect channels and took major accounts internal
- Reseller/Integrators could not justify expensive engineers and architects – became "order takers" for OEMs
- By 2006, OEMs realized costliness of internal sales staff: initiate layoffs and begin apology tour to channel partners
- Channel partners strike back: demand better margins in exchange for renewed trust

Independent Integrators on the Rise

- Often IT professionals impacted by downsizing (and some former vendor engineers/sales folk seeking to redeem their souls from misdeeds in their former roles within OEMs)
- Leading with "service codes of ethics" and "best practices" based on experience
- Some venturing into aftermarket service and support and used gear sales

The Latter Trend Underscores Another Significant Cost Issue

COST TO BUILD

- Raw Materials/Components
- Assembly Costs
- Value-Add Software
- Supply Chain Management
- Direct/Indirect Channels
- Distribution Costs
- Warranty and Maintenance
- RMI Churn
- Legal Fees and IP Protection
- Hedging

Hardware Costs Little, Software Lots

- Consumers frequently complain that embedded "value-add" software cost is largely invisible until year 2
- In years 3-5, you are effectively re-purchasing the entire array (which was declared end of life by its vendor 18 months after original purchase)
- Classic strategy…

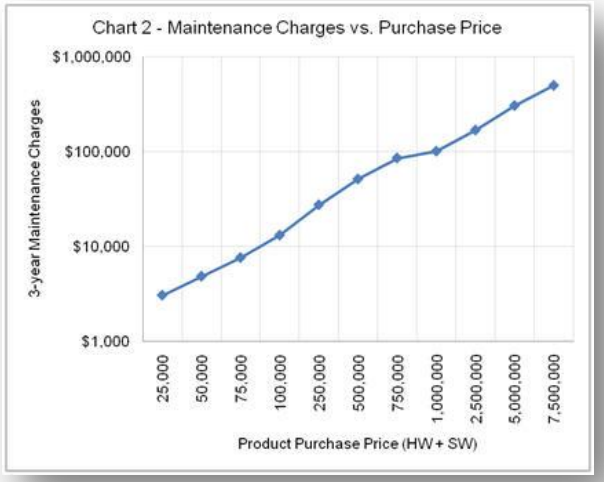

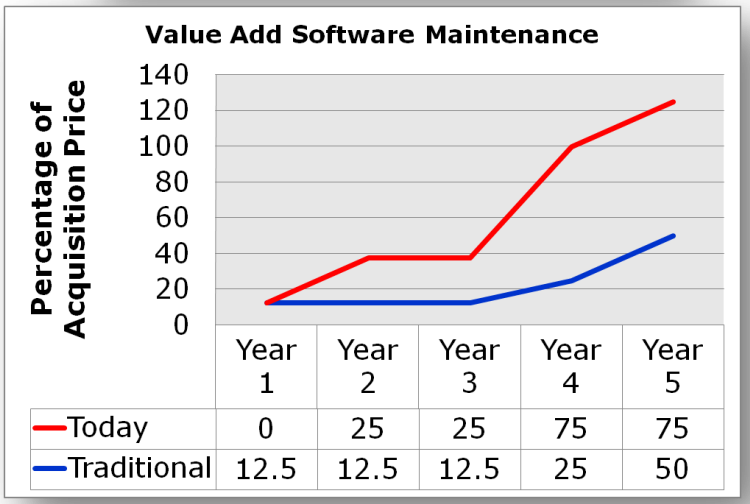

Razors and Razor Blades

- $750 million in R&D
- Three blades for closest shave possible
- Less skin pressure, fewer strokes for less skin irritation
- The best a man can get

- Latest in blade technology
- Breakthrough in comfort technology
- Latest innovation in personal grooming
- Micropulses for closest shave possible
- Five blade shaving technology

A Recent Storage Vendor Pitch

- Prominent NAS vendor offers array with lots of value-add software
  - Pitch: "We're confident that custo mers will use 50% less storage with [our st ag s we're willing to put our money wh mouth is with the $1 Million Challenge contest...[Our]sto technologies such as thi e-duplication, special RAI not an help realize unprecedente p e vings in virtual infrastructures..." or yo get 1M in products and services from the vendor.
  - Reality: Assume their gear doesn't perform. What is the value of $1M?

WHAT A DEAL!

1/5th of a NAS

Basic Rig
With SW
and 500TB of Disk
$5,039,884.30

Almost 3 Years Support (No NAS)

3 Year Maintenance Contract
$1,198,255.40

$1M Worth Much More in Used Equipment Market

- Late model hardware (usually re-certified by the vendor) at least 60% off sticker
- Upside:
  - Given reputable dealer, excellent service and support either from OEM or qualified third party support organization
  - Abundance of used equipment brands and models available
  - ASCDI does good job of policing up industry, provides forum for redress (ASCDI.org)

- Downside:
  - Not the "latest razor blade" technology
  - Some OEMs (NetApp for one) forbid software license transfer (limits residual value of products)
  - Not all aftermarket vendors reputable

3rd Party Maintenance & Support

- Not comfortable with this idea? Chances are you are already using it.
- Check out TSAnet.org
  - A few storage companies that use this vehicle to support integration in heterogeneous environments include…
  - While claiming not to be a 3rd party tech support organization, a substantial number of members use TSANet to deliver 1st, 2nd or 3rd tier support to customers, especially on older products

Using Qualified 3rd Parties for HW/SW Maintenance and Support

- Can shave up to 60% off annual maintenance costs
  - In the NetApp user community, Zerowait is an established brand: part high-availability storage engineering firm, part 3rd Party support firm for Netapp gear
  - Zerowait CEO Mike Linett says that just telling the vendor that you are looking at his company for alternative services will usually encourage NetApp to drop its prices for hw, sw and service substantially
- Leverage is a beautiful thing

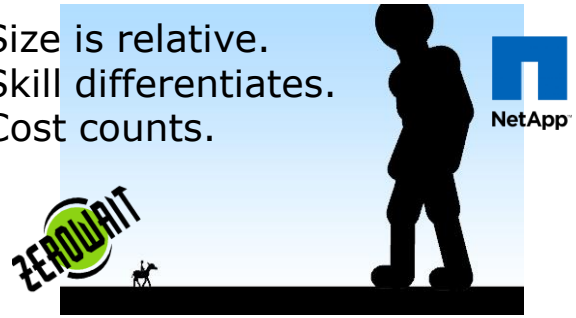

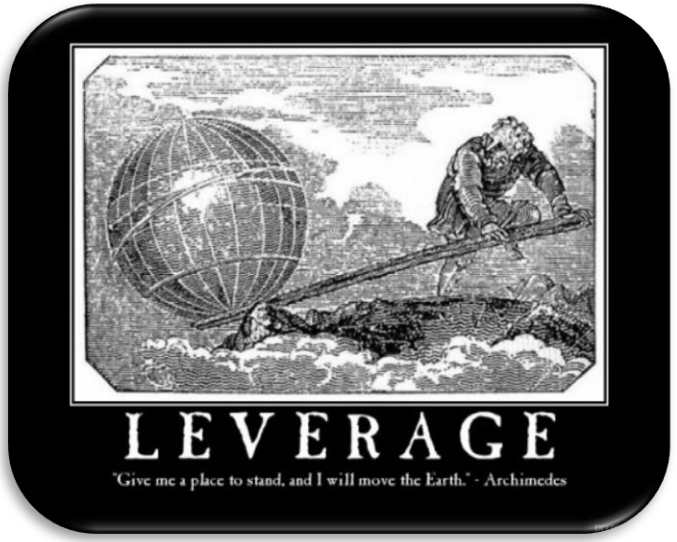

Summarizing

- Contemporary storage solutions reflect both "fixed" and "flexible" cost components
  - Fixed: Hardware is largely commoditized, channels add cost naturally
  - Flexible: Value-add software functionality, warranty & maintenance agreements
- Vendors really can't be villanized for price gouging: profit margins are actually quite thin
- No one I know pays MSRP!
- To contain costs, first we need to look at the **"flexible"** cost components

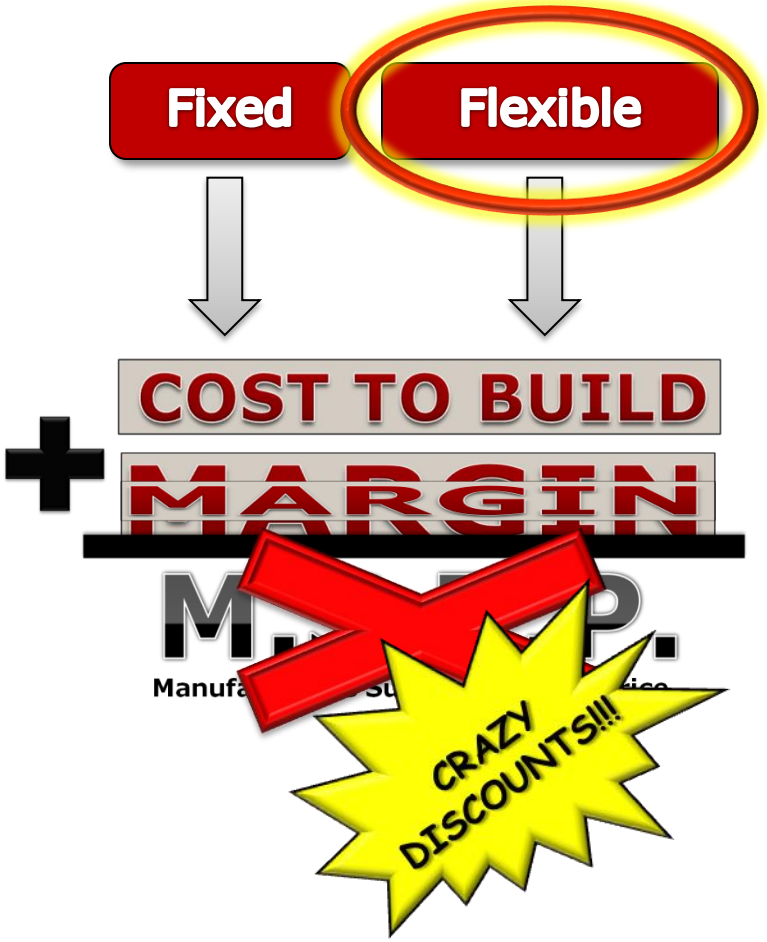

What Really Needs to Be Contained: Cost of Ownership

- Presentations like this wouldn't be complete without some helpful hints for applying what has been discussed in your world…
- I am not an analyst or a guru, but I believe that the longer term costs of misguided storage acquisitions and architecture are much more important than the tactical matter of getting a good price on gear.
- So here are my top ten pieces of advice, which, like me, you probably learned from watching television as a kid…

Ten Recommendations

1. Before you buy anything, understand what kind of data you are planning to store.

   Analysis should include:

   - Type of data (files, email, database, rich media, etc.)
   - Rate of growth
   - Access frequency
   - Concurrent access requirements
   - Sensitivity of data
   - Data protection requirements

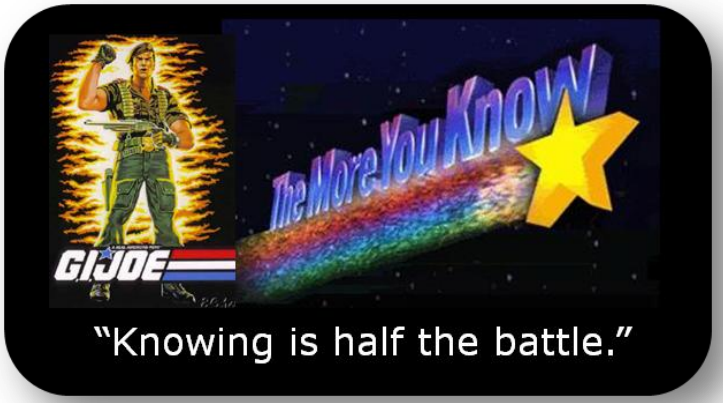

Ten Recommendations

2. Violence is the other half of the battle. To get a good price, you need leverage…
   - Eschew "love brands" – make ALL vendors compete for your business
   - Issue an RFI that specifies exactly what you need in terms of functionality and ESPECIALLY manageability
   - Consider used gear and tell the vendor that you are examining this option
   - Watch out for gotcha's
     - Scalability
     - Software transferability
     - Warranty and maintenance costs
     - Even "little things" like signed drives, capacity holdbacks, drive sparing ratios

Ten Recommendations

3. Trust is hard to come by these days. Verify via phone or on-site visit, not via *University of Google*
   - Independent analysts no longer that independent (if they ever were)
   - Bloggers have hidden relationships with vendors
   - Comments on forums, blogs and Twitter often conceal true identity of speaker (vendor ringers abound)
   - Trade press articles often slanted
   - Vendor gag orders impair free and factual exchange of performance and service quality data

Heads, I win.
Tales, you lose.

Ten Recommendations

4. Gadgets and gizmos don't necessarily translate into good storage.
   - Today's on-hardware features will likely become part of tomorrow's generic file systems (deduplication, for example)
   - Deconstruction and simplification are almost always better than "gadgetization."®

Is there a Tier 0 Flash RAM disk, as well?

Ten Recommendations

5. Treat disk like inventory. Don't maintain more than you need.
   - Disk capacity keeps improving, while cost per GB keeps falling: buy capacity when you need it – it is cheaper!
   - Thin provisioning and on-array tiering does not really improve capacity allocation, since you still need a large capacity buffer to insulate against margin calls: think about it.

C'mon. How many costumes does Ironman really need?

Ten Recommendations

6. Test everything. Almost no technology performs as described in the brochure under your workload.
   - Insist on proof-of-concept before you buy
   - Or get a test rig and test it yourself or using a trustworthy third party

Nothing works like vendor said. ***Hulk mad…***

Ten Recommendations

7. Find a cool sidekick. A competent solutions integrator can be a great asset.
   - Look for deep tech expertise and impeccable vendor neutrality
   - Performance-based contracting is always a good idea

Um.

Better.

Now we're talkin'.

Ten Recommendations

8. Everything fails. Value-add software-laden arrays fail more often. How will you recover?
   - Pin down the details of the vendor's disaster recovery story: what will help you predict and prevent avoidable disasters, and how will you recover data if the rig fails completely?

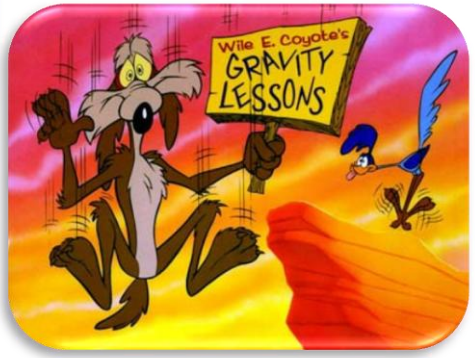

Ten Recommendations

9. You can't optimize infrastructure you can't see. Select a management method FIRST, then refuse any equipment that cannot be monitored and managed using your method.

I can't sense your storage utilization. How many element managers are you using anyway?

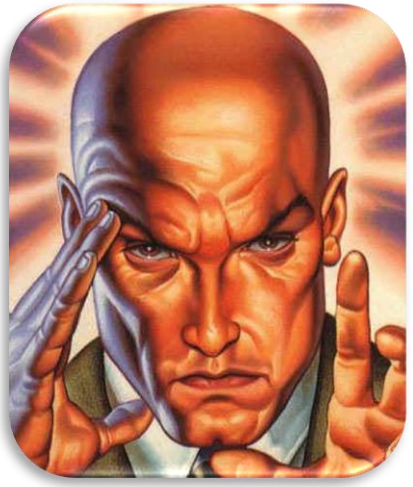

Number 10

10. Despite the many super powers evidenced in movies, television and pop culture today, you can make do with one: your common sense. Business requirements need to be aligned with technology capabilities today more than ever before. To make that happen, common sense must be your guide.
    - Buy what you need, not what the vendor wants to sell you.
    - Be circumspect about "too good to be true" solutions: one-stop-shop stovepipe arrays, FlashRAM drives, cloud storage…
    - Insist on standards-based technology: a lot of today's vendors won't be here in a year or two.






IT-$ENSE

Thank You.

- Questions?

- Feel free to reach out
  - Email: jtoigo@toigopartners.com
  - Twitter: @JonToigo
  - Blog: www.drunkendata.com
  - Websites: too many to list (pushing www.IT-sense.org today)
  - Blooks: www.makingitmatter.com, book.drplanning.org, www.datademocracy.com