Storage: Cost, Classification & Process

Anthony G. Mungal
Senior Consultant

10/8/13

Abstract

Storage continues to be the largest portion of the IT budget in many Fortune 500 companies. In fact, many of those companies have resorted to building simplified cost models to track the unit cost of managing a Terabyte per annum. While these costs reflect many of the obvious line items such as media, maintenance, leasing etc., much of what is offered as optimization techniques and best practices go unnoticed even with their ability to lower cost even more.

This session discusses the classification of storage necessary to satisfy the requirements of today's evolving workloads which run on "traditional", cloud, or hybrid storage and computing infrastructures. It examines some of the optimization techniques such as virtualization, de-duplication, virtual provisioning and others. It also looks at trends which indicate that economies of scale doesn't always scale and proposes a method for indicating the same. It carefully considers the impact of cloud versus non-cloud decisions on the overall cost model. Finally, it considers how capacity planning is to evolve in the presence of finer, more delineated metrics which facilitate all of the above.

Agenda

- Why trending is important.
- Architectural review of Oversubscription
  - Potential capex and opex savings opportunities
- Generalized Storage Cost Model
  - Determining "economies of scale"
- Identify and propose elemental, compound and derived metrics to enhance storage efficiencies
  - Performance and Capacity Planning

Some not so Obvious Trends

- Storage continues to be the largest portion of the IT budget in many Fortune 500 companies.
- Economies of scale doesn't always scale.
  - Scaling dependencies vary
- Many customers have experienced long term media cost trends with steadily decreasing halving times such that:[1]
  - For tape, it is exponential decreasing with about a 3 year halving time.
  - For disk, it is also exponentially decreasing with an 18 to 24 month halving time.
- The storage cost model can be misleading.
  - Cloud vs non-cloud considerations
- Storage capacity planning continues to evolve in the presence of finer, more-delineated metrics.

Storage Stakeholders

- Virtualized environments …private cloud
  - Request for VM includes some amount of storage
  - E.g. 50 GB to 70GB per VM for the OS; +4-32 GB for swap space
  - Data: ??
- Explicit storage request via email from server/application team specifying a quantity of storage
  - High number of GB…getting to be in TB amounts.
- Above defines "allocation"; a subset of which defines "growth"
- Storage funding
  - Capital cost (project level or enterprise wide)

Process - Shortcomings

- Need for specification of performance SLOs, availability, or retention criteria.
  - Will help with FA over-commitment issues
    - E.g. tier 1 FAs to have higher FA:pool ratio <u>or</u> all FA:single pool monitoring.
- No tracking of unused space within the Databases
  - Significant % of space allocated but unused.
- Overall management of allocated space is poor
  - Finding: lacks process…not so much tools!
- No visibility into the VMs
- Other

Storage Classification

- Total Raw Storage
  - Defined as the sum of all storage prior to the necessary formatting and protection schemes being applied to render it "usable".
- Total Configured Storage
  - The subset of raw storage which has been made "usable" and from which users requests can be fulfilled.
- Total Allocated Storage
  - The subset of configured storage which has been allocated to the various servers.
- Total Used Storage
  - The subset of allocated storage that is actually used by server(s).

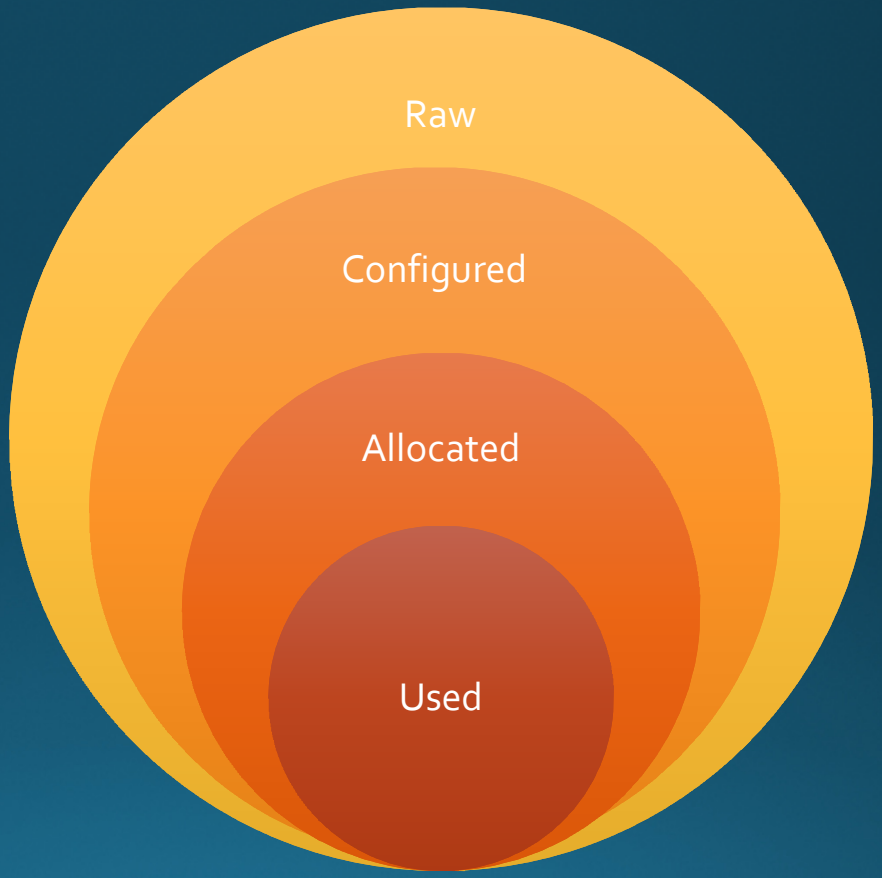

Thin/Virtual Provisioning

- Some amount of configured storage is logically grouped to form a "thin pool" which has the following properties:
  - Physical capacity of *x* GB.
  - Logical capacity of *kx* GB such that *k* is usually around 1.2*
  - Pool is seen as k times larger than it really is.
  - *k* is known as the "oversubscription" factor and is a dynamic metric.
    - k=1 , no oversubscription (like the "non thin provisioning" times)
- User allocation requests are fulfilled from this pool
- Oversubscription allows more storage to be allocated than is currently configured.
- Adding more configured storage can be done at a later time
  - Cost deferment - CaPex, OpEx impact
  - Does not hold for rapidly increasing values of *k* once *k* > 1.0

* as dictated by the storage demand growth rate; for lower rates, k can be higher, and for high, aggressive rates, k should be lower.

Thin/Virtual Provisioning - Changing The Rules

- Allocating what the user requests … but not really!
  - Tell the user that you allocated the amount of storage they requested.
  - Storage subsystem honours the request but makes available to the user the storage only as it gets "used".
  - Most users don't use all of what they ask for, and hence the reason thin provisioning is popular and successful.
  - Used ≤ Configured

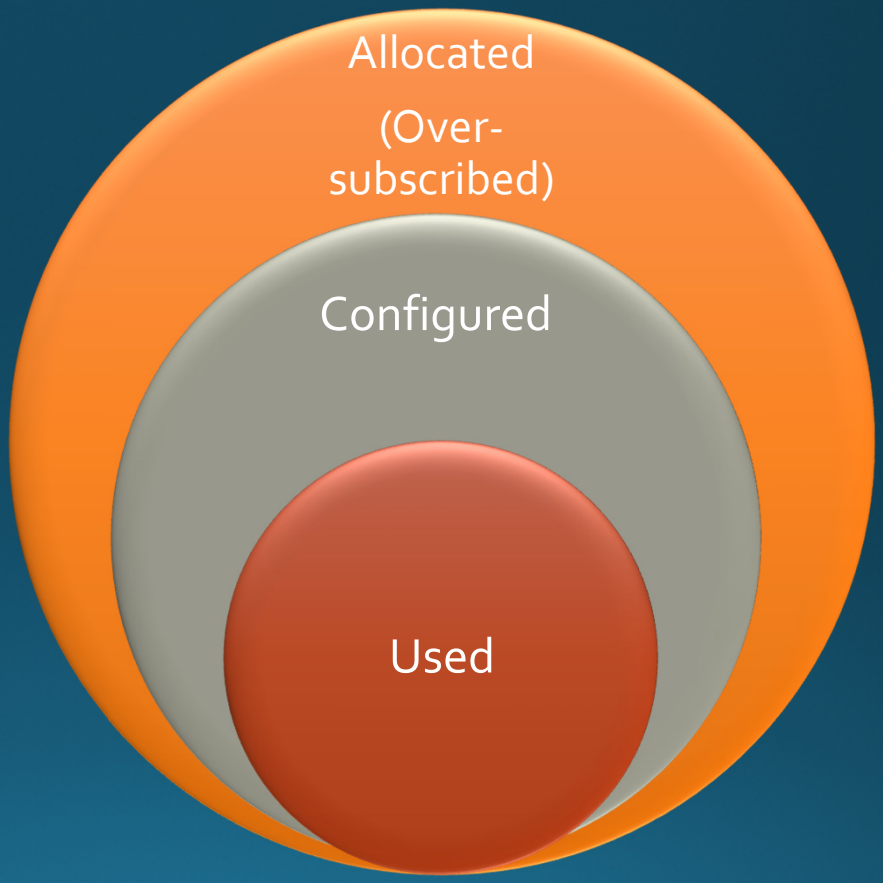

Understanding Oversubscription and Growth

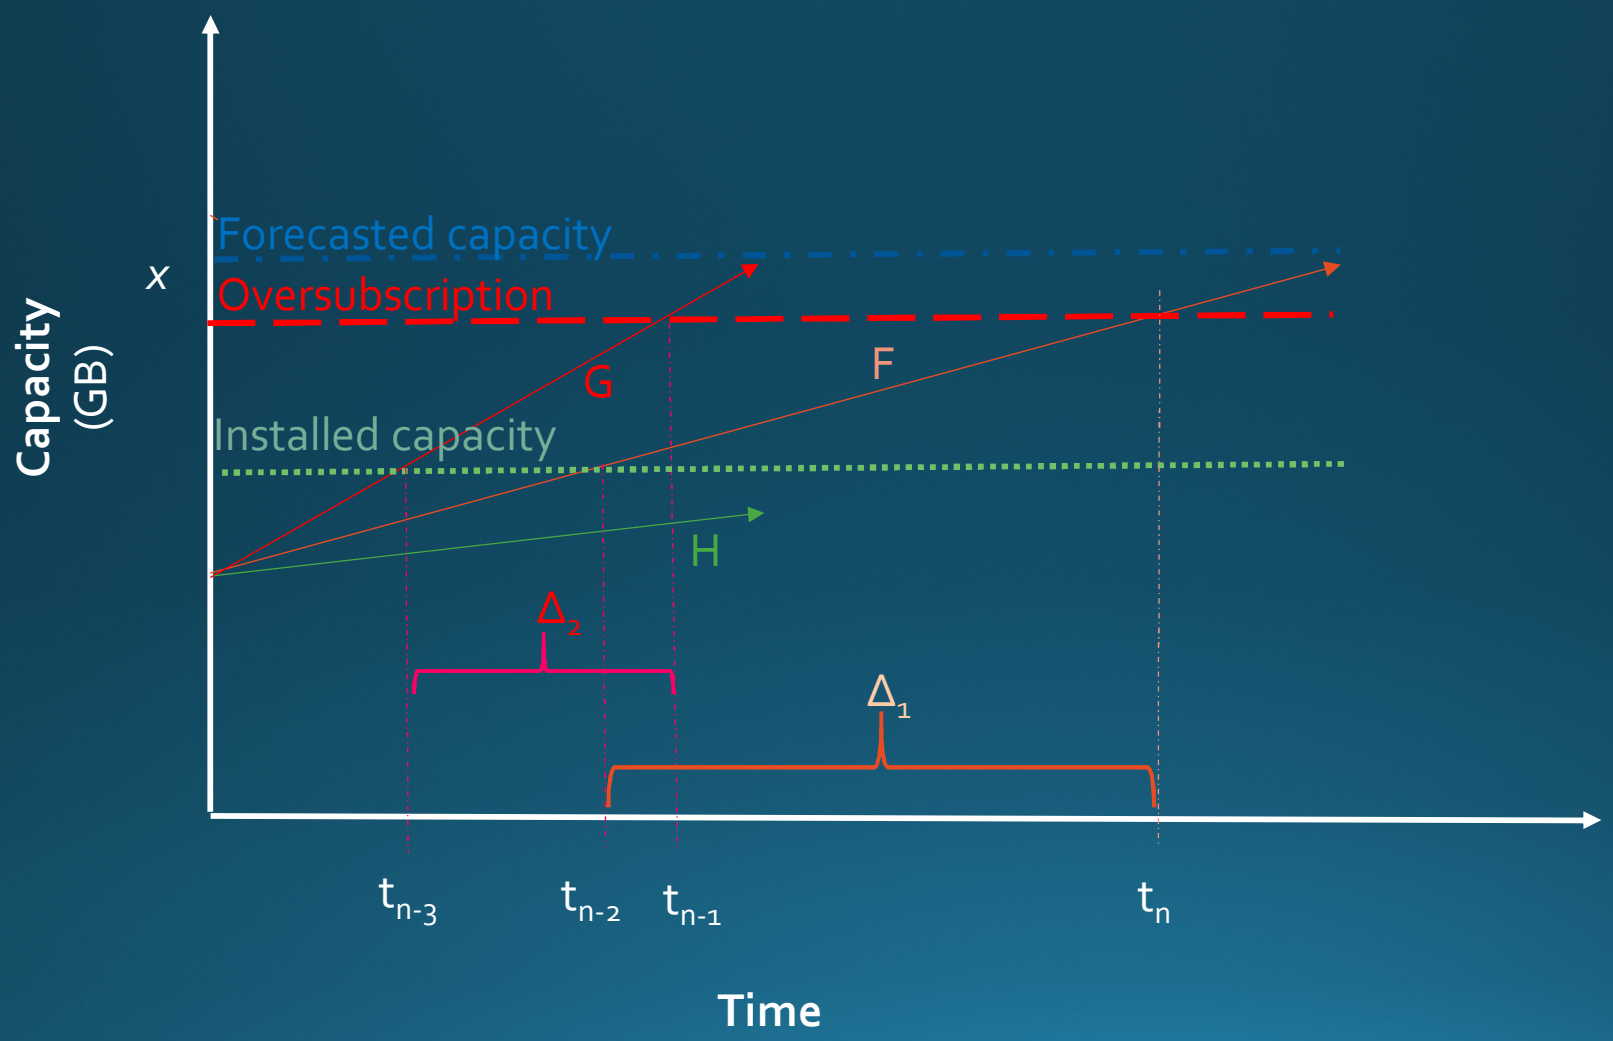

Cost Model Introduction

- Near exponential data growth heightens the need for understanding cost of storage (past, present and future via projection/forecasting).
- Basic equation:
  - Cost = vendor hardware & software cost + customer "operational" cost.
- Operational cost differs from one environment to the next.
- Cost Elements include:
  - Servers and supporting hardware
  - maintenance
  - software licences
  - floor space
  - utilities
  - labour
  - other

Generalized Non-Cloud Cost Model: Disk & Tape

CRITERIA	DISK		TAPE		Comments
	$/TB/yr	% of total	$/TB/yr	% of total	
Subsystem & Media					Modest % of total?
Other CaPex (filesystem servers, SAN, etc.)					
Maintenance & licences					Increasing % of total?
Facilities (space, power, etc.)					Increasing % of total?
Admin ($/FTE)					Increasing % of total?
Network access cost					
Optimizations (e.g, FAST/VP, oversubscription, de-dup, etc.)					Decreases $/TB/yr
TOTAL	$x/TB/yr		$y/TB/yr		

Economies of scale is derived by determining which costs scale proportionately with size; which remain fixed ; which scale only weakly with size.

Some Benchmark Behaviours

- Unabated storage growth (high CAGR …to exponential)
- Media cost/byte is lowered as economies of scale is exploited (infrastructure growth)
- Media capital cost, maintenance, and facilities usually scale linearly with size.
- Some licence costs remain fixed; some others will scale weakly.

Some Benchmark Behaviours … cont'd

- Some of the supporting infrastructure scale weakly with increased size (also infrastructure may be fixed in cost until a certain threshold i.e. capacity is reached).
- Disk cost is determined by amount of disk deployed whereas for tape it is determined by amount of data stored; better normalization required. For virtual tape, this argument is mooted somewhat.

Changes to the Cost metric

OLDER SUBSYSTEMS

NEWER SUBSYSTEMS

	COST		COST	IOPs
Enterprise	$k/GB	SSD	$x/GB	70% ?
Mid-tier	$(k-n)/GB	FC	$y/GB	20%?
		SATA	$z/GB	10%?

Composite cost:
$(k + (k-n)) / GB

Composite cost:
$(x + y + z)/ GB

KEY POINTS:

1) Comparison: (x + y + z) ≤ (k + (k-n))
2) Newer subsystems can do significantly more IOPs through the use of multi-tiering & FAST/VP while at a lower cost point thereby reducing cost per I/O dramatically.
3) Cost per I/O becomes a very effective metric.

Capacity Planning

- High Level Architecture
- Goals and Objectives

Capacity Planning – Main Premise

- More effective use of resources
- Defines infrastructure limits (ceiling function).
- Lowering of the "effective" cost of the infrastructure …through higher utilization levels and improved service delivery.
  - lowering OpEx.
  - CapEx avoidance.
- Facilitates constant evaluation of the infrastructure against sole or shared alternatives (inhouse, and private, public, hybrid cloud)
- Matching of business growth with current and future technologies.
- Anticipatory versus Reactive modes of operation.

Overview

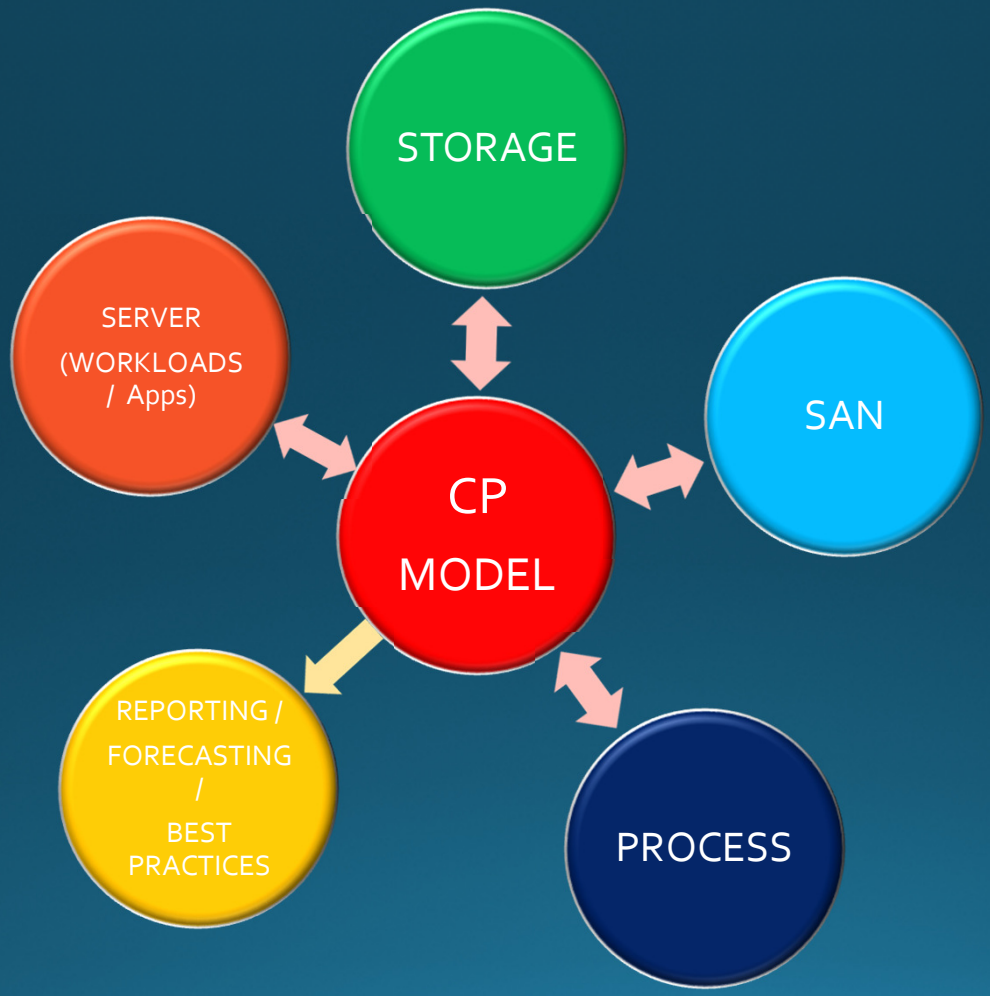

The Pareto Principle

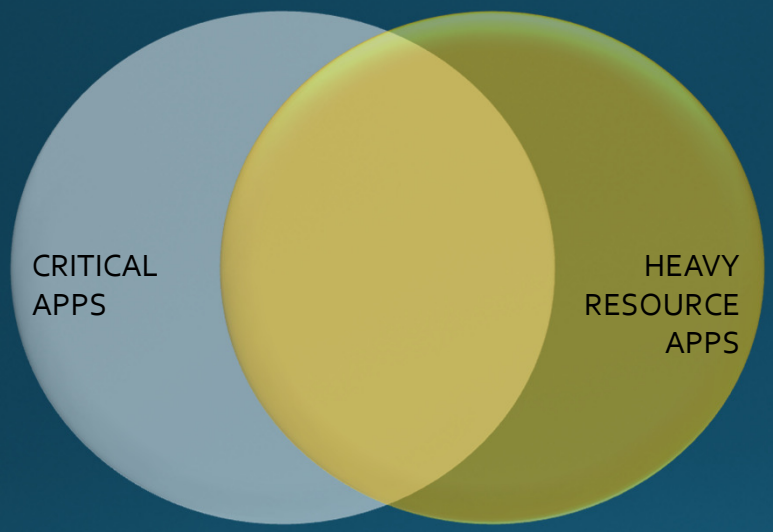

- *k% of workloads consume (100 – k)% of the resources.*
  - E.g. ≈ 20% of the apps are critical…accounts for ≈ 80% of the revenue…consumes ≈80% of the resources/manpower.
- The heavy resource apps are usually a small % of the app population.
- Specialized capacity planning for these two categories of apps will render big payback (worthwhile effort).

Capacity Planning High Level Architecture

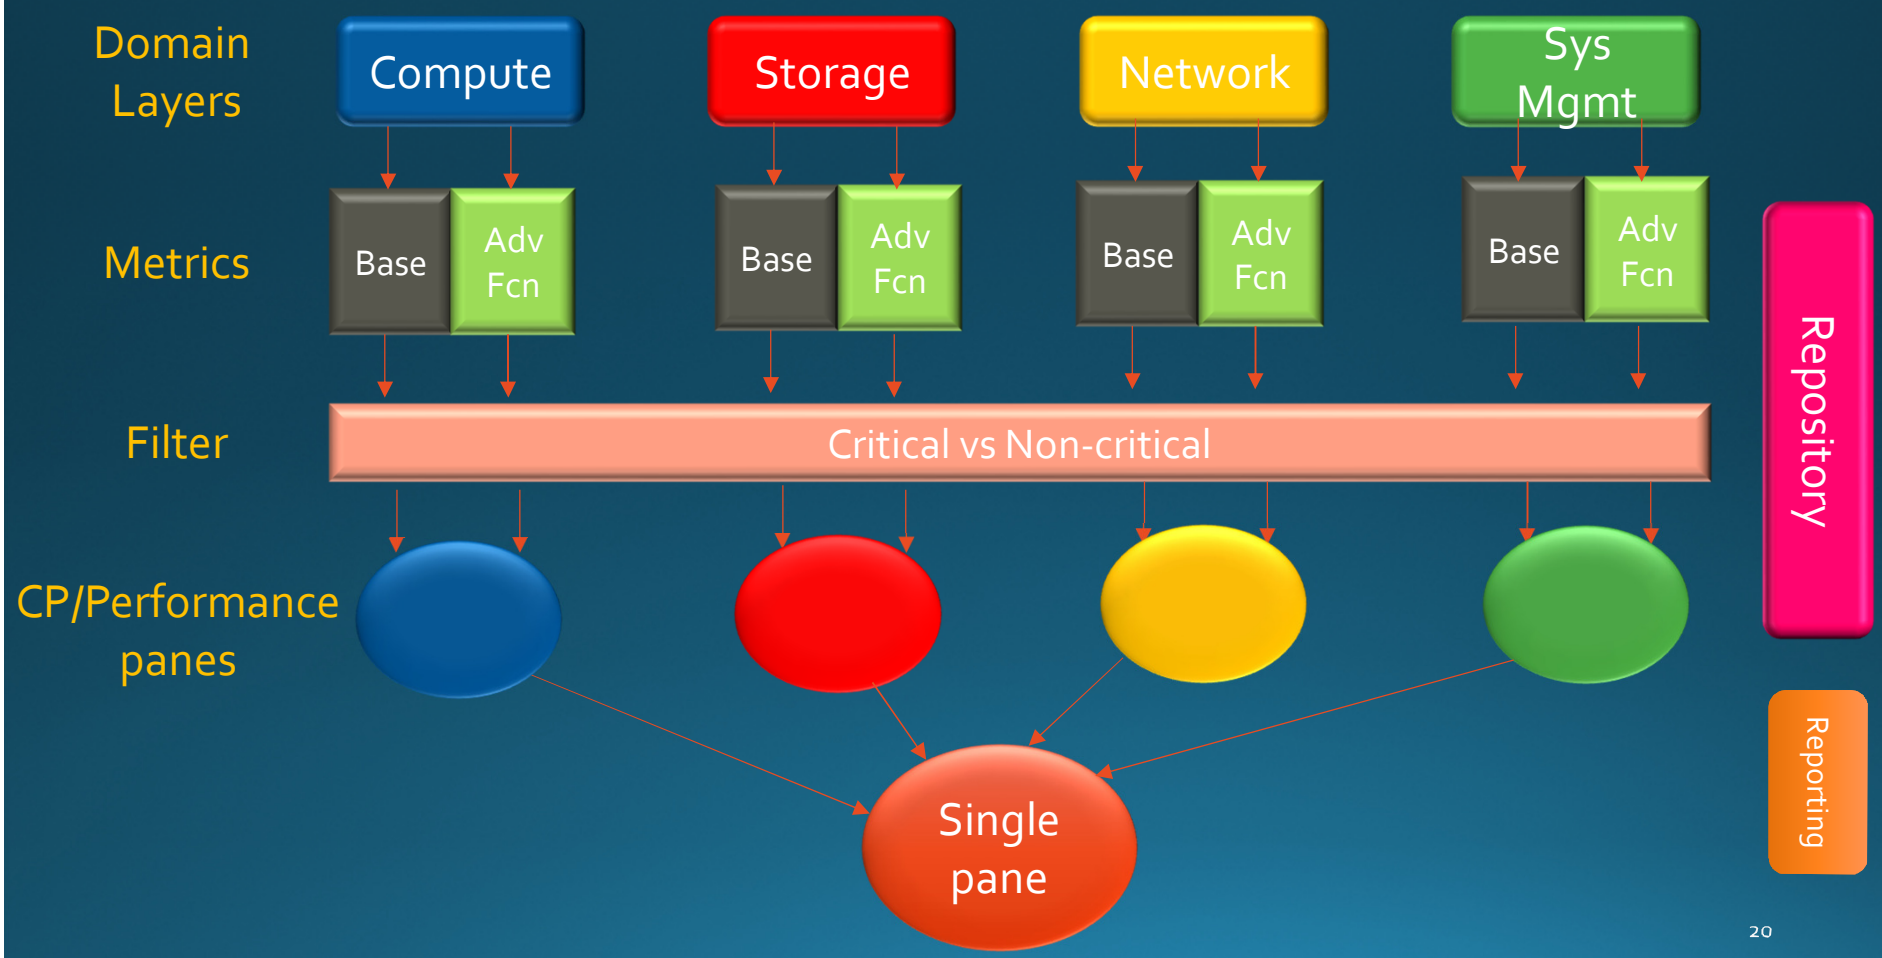

Process

- Manual
  - Pull process oriented
  - Data capture and analysis of KPIs
- Automatic
  - Push process oriented (e.g. alerting)
  - Automatic monitoring of critical components.
- Actionable vs non-actionable

Process: Static & Dynamic Thresholding

- Static Thresholds
  - Very well understood and used.
  - Thresholds defined as a limit against which actions can be taken. e.g device busy > 45%, I/O rate <1000, etc.
  - Alerts can be purely informatory or actionable (immediate or deferred).
- Dynamic Thresholds
  - Powerful, sophisticated method based on statistical measures which convey trend and behaviour tendencies.
  - Thresholds set against varying, statistical numerical limits which are based on historical trends:
    - e.g. "Tier 1 allocations > 2σ more than usual for this day at this time"
  - Application of Multivariate Adaptive Statistical Filtering (MASF)

Reporting & Forecasting

- Reporting
  - Frequency of reports (daily, monthly etc.)
  - Critical workload vs non-critical reporting
  - Reporting interface
- Forecasting (growth prediction)
  - Trending (intervals?)
  - Workload level
  - Resource level
    - Component level
- Best practices

Storage Capacity Planning – Goals and Objectives

- Capture, analysis and usage of key storage performance metrics
- Faster problem resolution
  - diagnostic procedures which utilize timely information that is easy to understand and convincing to the parties involved.
- Application capacity trending and forecasting
- Lowering of the "effective" cost of SAN/Storage
  - higher, more optimized utilization levels and improved service delivery.

Data Sources

- "Tools slide" …later

Storage Optimization/Efficiency – Foundation Architecture

- Addresses:
  - Cost
  - Performance
  - Scalability

What's next ?

- "peel back the onion"
- Base metrics vs Advanced function metrics
  - Public cloud considerations
- Examine critical vs non-critical workload needs
  - Metrics required
- Data Sources
- Modelling & Forecasting requirements
- Other

Biography

Anthony G. Mungal is an independent consultant, author and speaker on many current IT topics. He is a 35+ year veteran of the IT industry and has earned numerous distinction and special recognition awards with companies like StorageTek, the Amdahl Corporation and the EMC Corporation while holding sales, marketing and technical positions which included: Consulting Systems Engineer, SE Consultant, Senior Product Manager, Large Systems Account Specialist, Regional Systems Engineer, Regional Product Specialist, National Systems Engineering Manager, DASD Product Manager, Senior Systems Programmer, Lead Business Systems Analyst, and Senior Programmer.

He is a graduate of the University of Toronto double majoring with honours in both Mathematics and Computer Science. He has participated on many IT discussion panels, co-authored three books and published numerous papers on Processor Performance, Memory Management, I/O Subsystems Configuration & Performance, Storage Management, IT Infrastructure and Architecture Design & Implementation, and other related topics which he has presented at major technical forums such as the Computer Measurement Group (National and Regional Conferences), CMG International Conferences (Australia, South Africa & the UK), SHARE, GUIDE and an assortment of other local IT related user groups and conferences. He is an active member, and holds executive positions with the Computer Measurement Group (CMG), CMG Canada, and Florida CMG (FLCMG). He is also an active member of both the Association for Computing Machinery (ACM), and the Institute for Electronic and Electrical Engineers (IEEE). He can be contacted at amungal@acm.org

Thank you