



Education

Best Practices in Managing Virtualized Environments

Mike Matchett, Akorri
10 April 2008

SNIA Legal Notice

- 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.

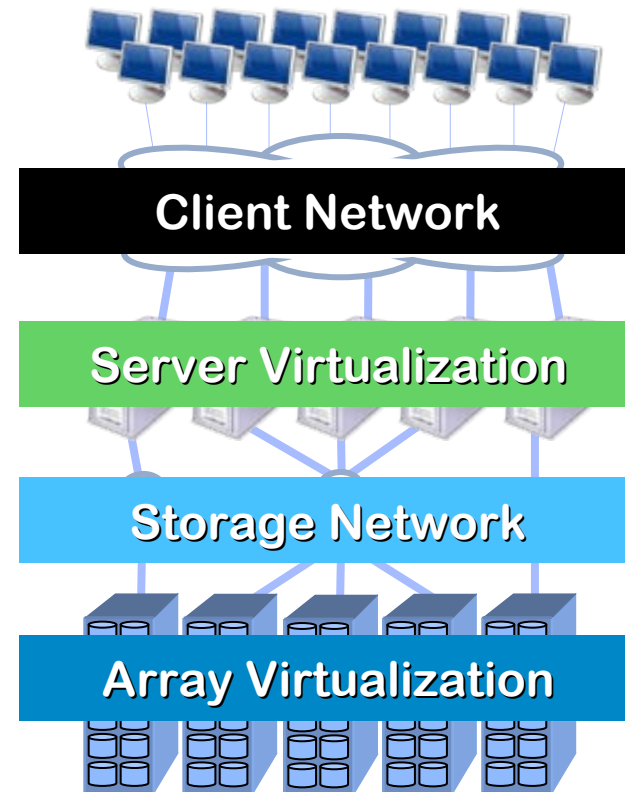
➤ Best Practices in Managing Virtualized Environments

- ◆ Today, the data center environments are getting increasingly complex with virtualization at all layers of IT stack, including network, server, SAN and storage. IT professionals are often challenged in diagnosing application performance issues, optimizing infrastructure resource utilization, and planning for future changes. The best practices for managing complex data center environments include cross domain management orientation, watch the infrastructure response time for cross-domain performance, look for application contention and contention-based latency in the storage layer, best fit analysis of workloads to storage resources, and work toward infrastructure performance SLAs. The key requirements for the new-breed management software include agent-less discovery and SMI-S support.

➤ Tremendous Benefits

- ◆ Pooling of resources
- ◆ Rapidly deploy new applications
- ◆ Increase resource utilization
- ◆ Over-subscribe resources
- ◆ Lower acquisition cost and TCO

➤ Traditional system management practices may no longer work



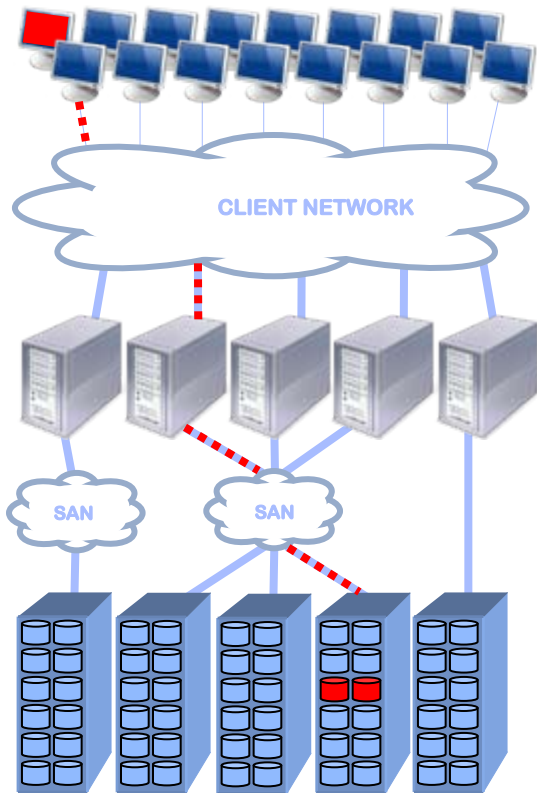
What's "Real" about Virtualization?

- Like the Emperor's new (virtualized) clothes –
 - ◆ A logical interface presenting a normalized "resource" that isn't "all there"
 - ◆ Built over physical and other virtual layers that may not look at all like the presented logical resource
- We are concerned today with the two major IT virtualization initiatives
 - ◆ Storage Virtualization
 - ◆ Server Virtualization
 - ◆ (and the combination of the two!)




**Check out SNIA Tutorial:
Virtualization I- What, Why,
Where, and How**

Physical Infrastructure Model



Virtual Infrastructure Model




- Managing through Virtualization is Challenging
 - ◆ Diagnosing Performance Problems
 - ◆ Optimizing Resource Utilization
 - ◆ Planning for Future Changes

<u>Example Virtualization Feature</u>	<u>“New” Management Challenge</u>
Clients can both Reserve and Share Resource Capacity	Resource Performance still Degrades Non-linearly with Load
Increased Resource Utilization	Optimal Resource Deployment
Easy to add to the Resource Pool	Predicting if the next “Client” fits

The Bottom Line...

- Applications share resources
- Poor performance is caused by:
 - ◆ Hard-to-find I/O bottlenecks and resource contention
 - ◆ Mis-alignment between layers of virtualization
 - ◆ Under-provisioning shared resources
- Over-provisioning of shared resources as insurance negates ROI
- Inhibitors to success
 - ◆ Virtualized data center complexity
 - ◆ Lack of cross-domain management
 - ◆ Lack of cross-domain communication



- Solving Old Problems in a New Environment
- Recommended Best Practices -
 1. Cross Domain Analysis
 2. Adopt an Application View of Performance
 3. Finding Shared Resource Contention
 4. “Effective Capacity” Management
 5. Model-based Optimization and Planning

1. Cross Domain Analysis

Virtualization Management is “Cross-Domain” -

➤ Create a Cross-Domain Baseline (discover and collect)

- ◆ Need data from multiple layers (app, server, storage, phys and virt)
- ◆ Aim for agent-less and “on-line”
- ◆ Standards like SMI-S are essential for heterogeneous environments

➤ Build in Automation for Scalability

- ◆ Virtualized environments tend to sprawl
- ◆ And they can change dynamically

➤ Check Configuration First

- ◆ Don't optimize or “plan from a baseline” from a poorly configured system
- ◆ Checklist vendor configuration best practices



Check out **SNIA Tutorial:**
Storage Management: SMI-S to
Management Frameworks

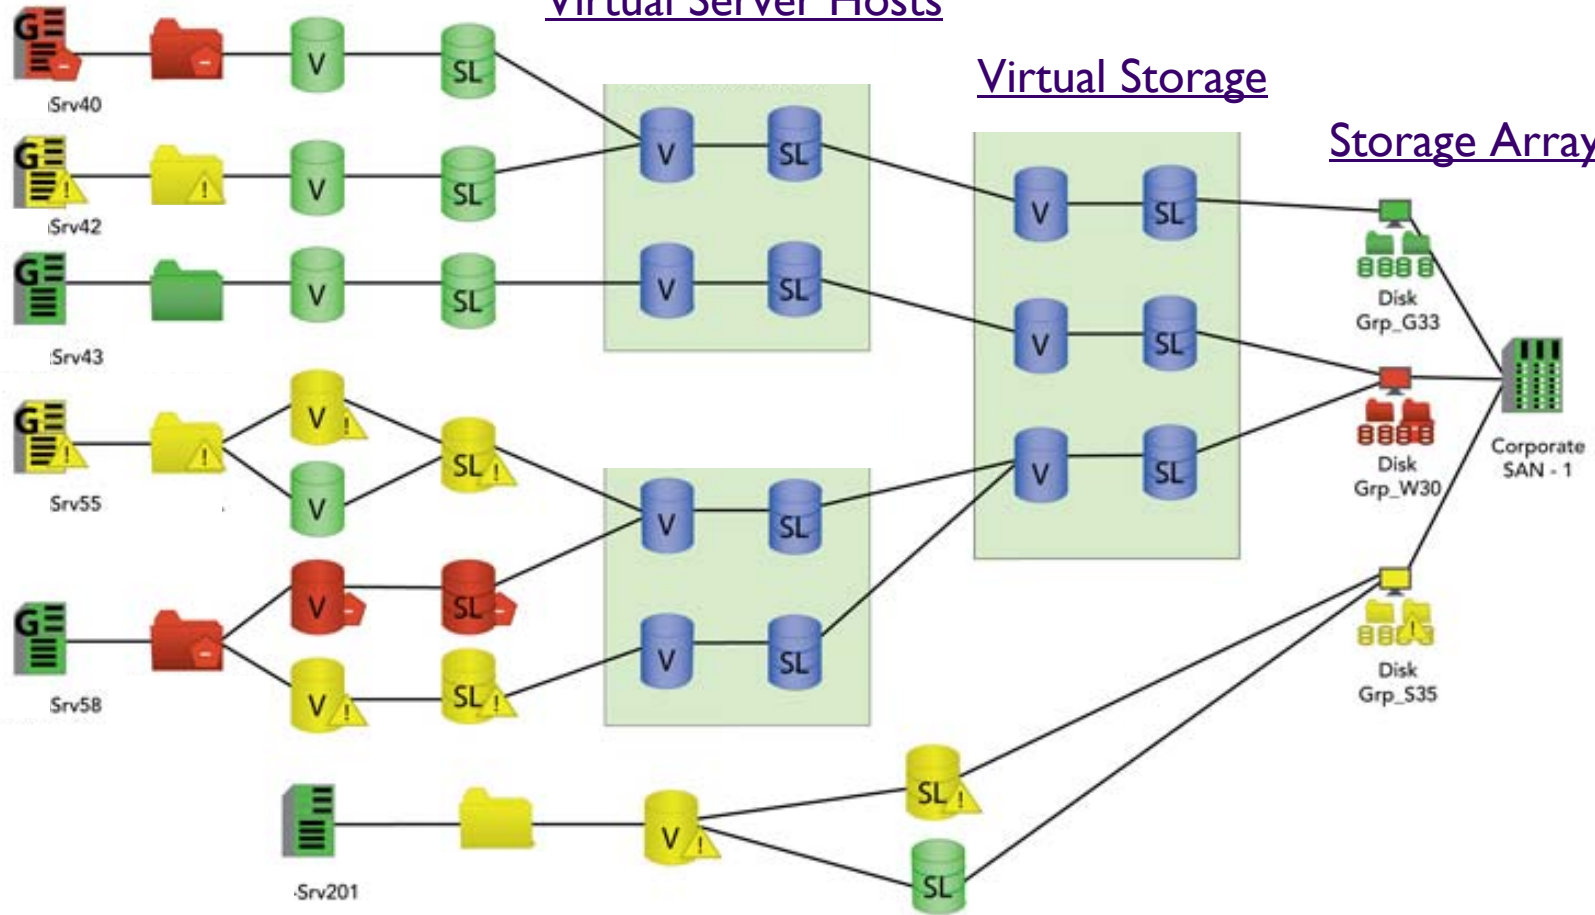
I/O Paths Through Virtualization

Applications and Servers

Virtual Server Hosts

Virtual Storage

Storage Arrays



2. Adopt an Application View of Performance

The Customer is Always Right –

➤ Application Infrastructure Performance

- ◆ How much time an IOP takes to complete from the application point of view (Response Time)

➤ Manage to this “Service” Performance

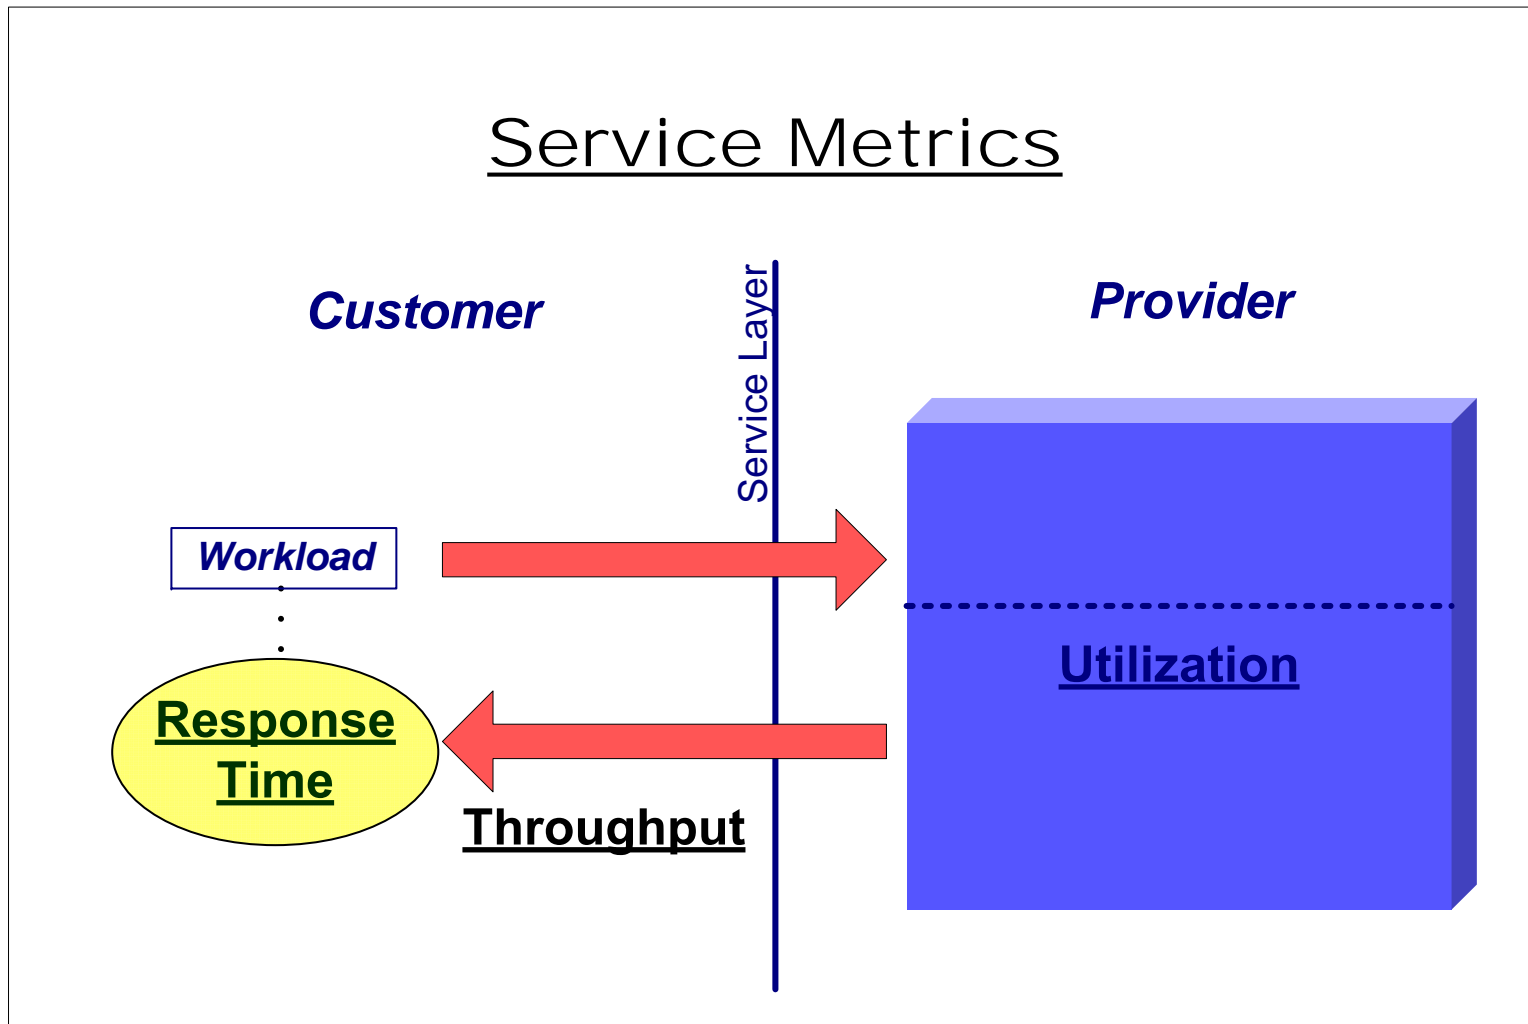
- ◆ Element utilizations are interesting, but service performance is the goal

➤ Look for Abnormal “Service” Behavior

- ◆ Not just default rule-of-thumb thresholds on utilizations



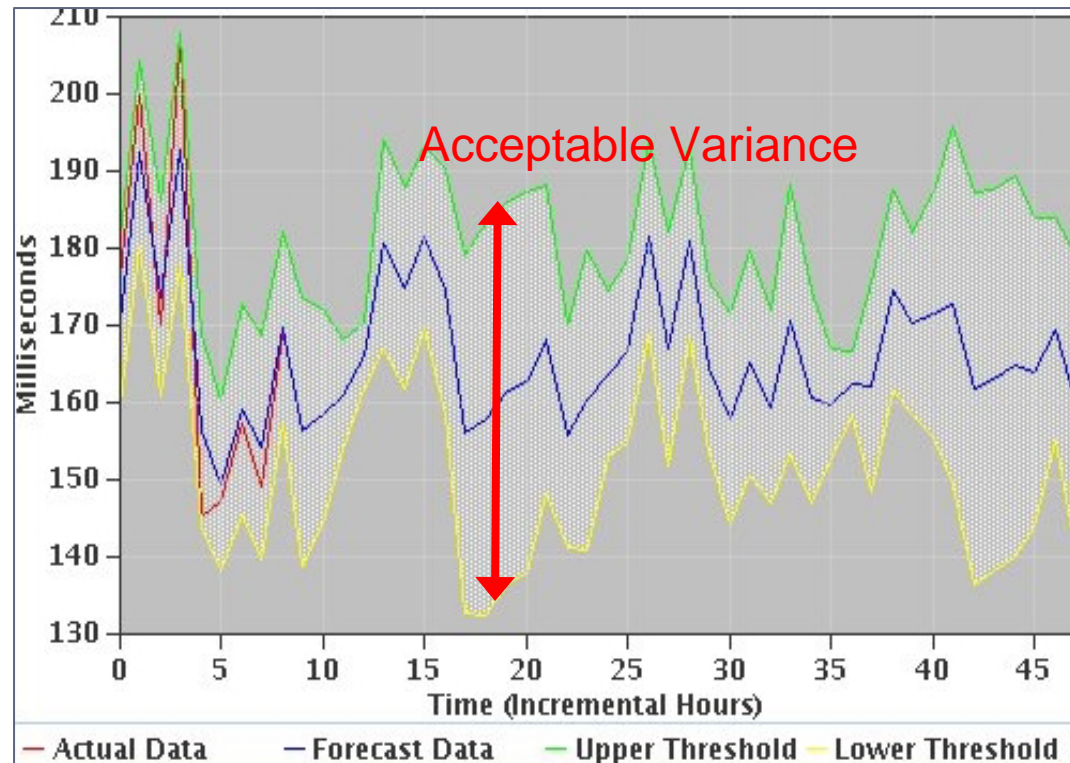
**Check out SNIA Tutorial:
Storage Performance 101**



Look for Abnormal Behavior

- Check for Abnormal Behavior
 - ◆ Calculate “Normal” baseline
 - › A common statistical analysis of variance of performance over time
 - ◆ Compare new data to baseline

- Shared Resources tend to average out peaks that will show in dedicated resources
 - ◆ Helps Justify Virtualization



3. Find Shared Resource Contention

Stepping Through a Virtual Looking Glass -

➤ Need to Map through Virtualization Layers


- ◆ Map relationships at every level
- ◆ Exponential problem with server virtualization over storage virtualization
- ◆ Sum up the loads from every client that shares each resource

➤ Quantify Application Contention due to Sharing

- ◆ Calculate performance impact back to each application

➤ Root Cause is mostly figuring out “What’s Changed” when Capacity runs out


- ◆ If Load, was it aberrant behavior or growth?
- ◆ If Configuration, does it violate policy or show thrashing?
- ◆ If Contention, who is new to the pool?



- **Cross Domain visibility is naturally “foggy”**
 - ◆ Domain specific management has limited view
 - ◆ Virtualization makes it harder
- **Management requires the end-to-end picture**

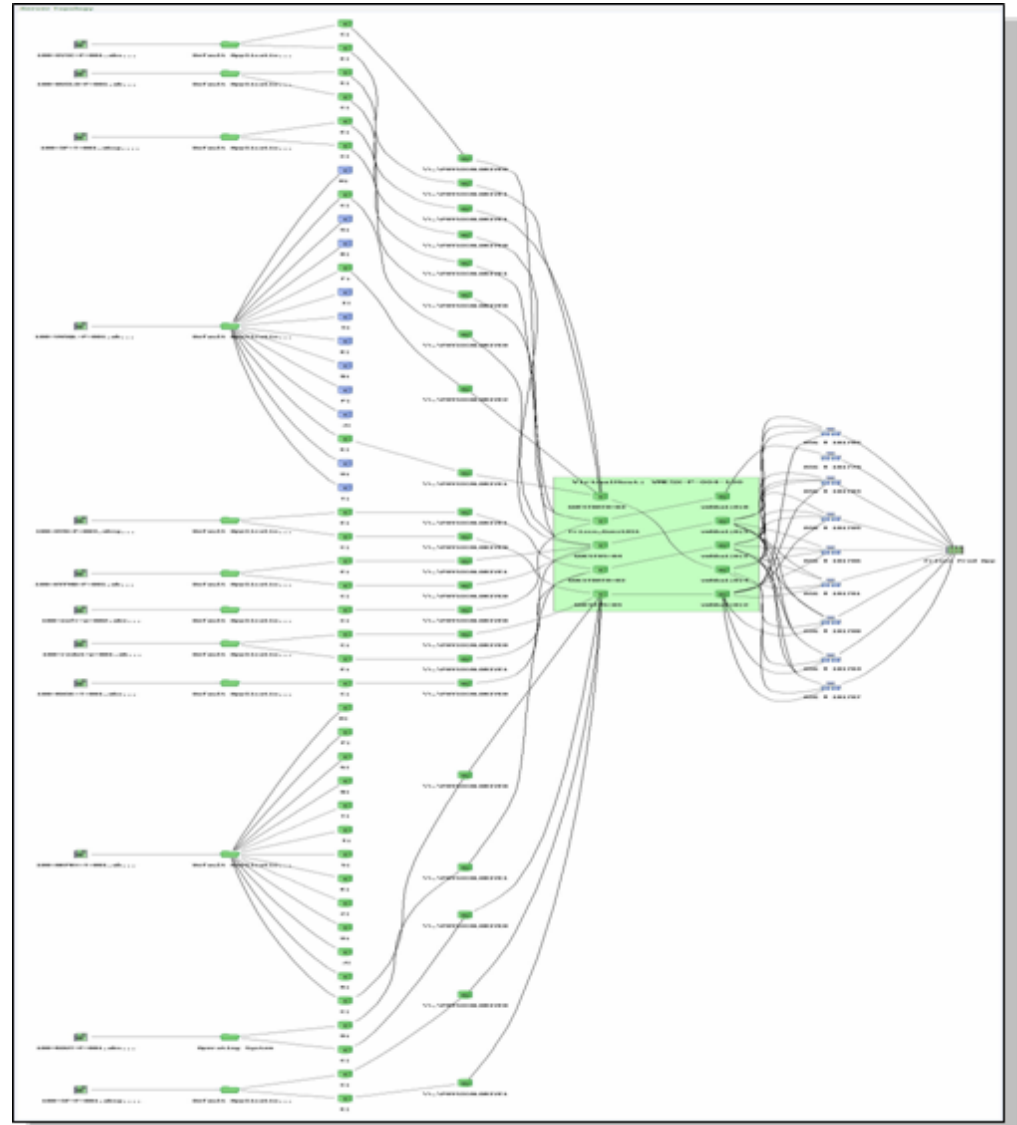
Cross-Domain: A Single VM Example

➤ Note the long data path from application to array...




Cross-Domain: A Virtualized Host Server Example

- Need a map through all the “indirection”
 - ◆ Sharing can be dynamic – maps must be too
- A common map helps different domain admins communicate



Cross-Domain: Server and Storage Virtualization Example

➤ Note the explosion in data paths!



4. “Effective Capacity” Management


Capacity Management Isn’t Just “Enough GBs”

- Storage has both space and time constraints (server folk have it easy!)
- Manage to the total “Effective Capacity”
 - ◆ Maximum utilization that gives good performance
 - ◆ Not to total actual utilization (aka “saturation”)



**Check out SNIA Tutorial:
Storage Virtualization II –
Effective Use of Virtualization**

Effective Capacity Is less than 100% SNIA



- Max. Capacity happens when system is 100% physically utilized
 - ◆ Response Time is technically infinite
- Service Levels determine acceptable performance (and optimal throughput)
- 100% Effective Capacity can be defined by that optimal throughput

5. Model-based Optimization and Planning **SNIA**

Move Towards a Real-Time Datacenter -

➤ Constantly Increase Operational Efficiency

- ◆ Most working infrastructure is sub-optimized
 - › Dedicated resources
 - › “If it ain’t broke, don’t fix it” attitudes (or capabilities)
- ◆ However, when everything is shared, everyone goes down together...

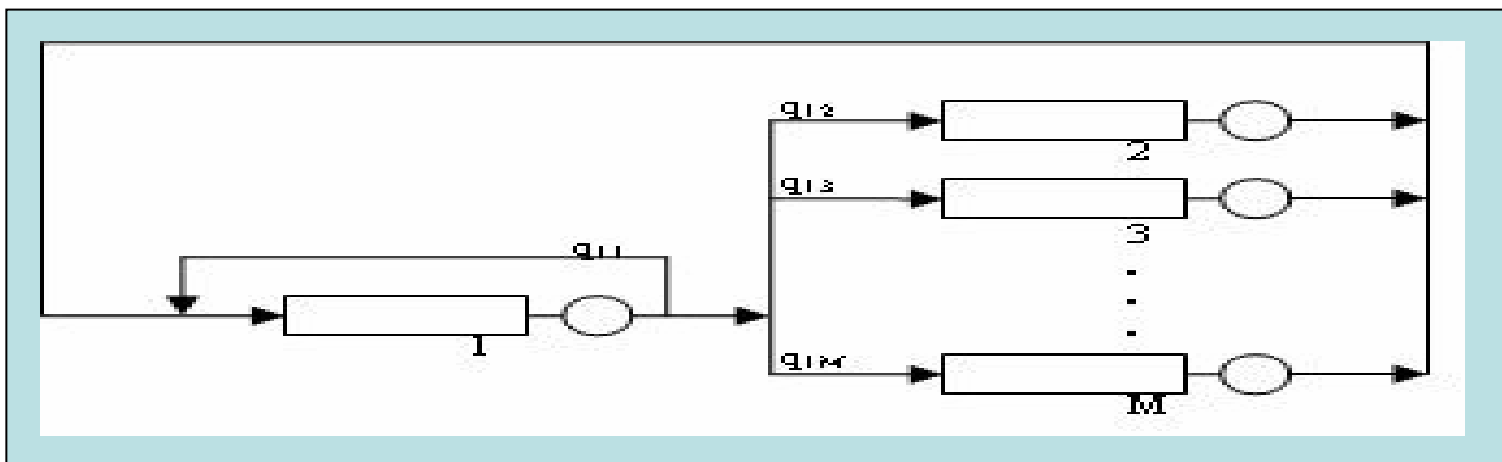
➤ Real-time Capacity Planning

- ◆ Utilizations can be related to Response Time through Queuing Theory
- ◆ Need to predict performance degradation under future application load changes
- ◆ Need to predict performance improvements from possible architectural/technology changes

[My prediction – The Planning “Cycle” over time will approach real-time...]

Queuing Theory to The Rescue...

- **Queuing Models create Response Time curves**
 - ◆ Based on established mathematics (Buzen, et.al – see www.cmg.org)
 - ◆ Useful analytically (historically) as well as predictively
 - ◆ For a simple example think of a check-out line at the grocery store
- **Complex Queuing Network Models can represent nested and virtualized IT domains**
 - ◆ Advanced cross-domain solutions model IT virtualization



➤ In Summary -

1. Cross Domain Analysis
Need to look across IT domains
2. Adopt an Application View of Performance
Manage to service levels
3. Finding Shared Resource Contention
Virtualization is about sharing, and that's often the problem
4. “Effective Capacity” Management
Shared resources still obey the laws of physics
5. Model-based Optimization and Planning
Leverage Prediction to Improve your Future

- Please send any questions or comments on this presentation to SNIA: trackvirtualization@snia.org

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

- SNIA Education Committee

**Rich Corley
Peter Lauterbach
Tom Joyce
Kevin Faulkner**

**Rob Strechay
Brian Schofer
Lisa Crewe**