DIFFERENTIAL VIRTUALIZATION FOR LARGE-SCALE SYSTEM MODELING

Jason Koppe
CSIA Master's Thesis

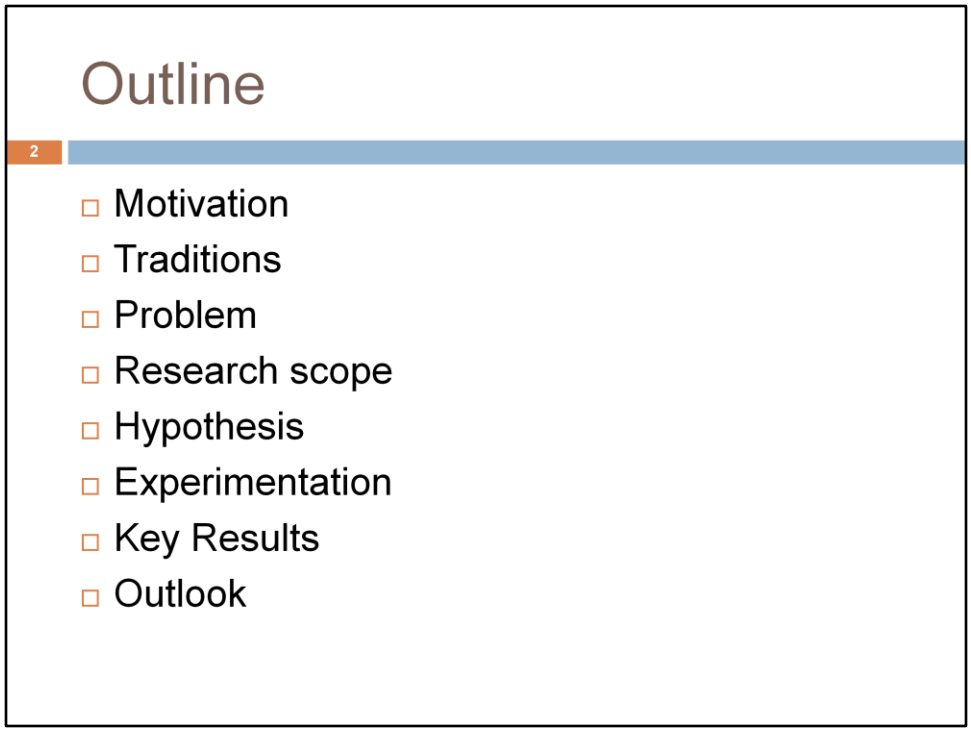

Motivation – how did we get to researching this area?
Problem – system modeling isn't scalable to a large number of preservable models with regards to process performance and management
Research Scope – to what does this apply?
Hypothesis – a new modeling process (encompassing new administrator and user processes)
Experimentation – what can we do to understand whether this hypothetical environment scales? How can we make such an environment manageable?
Results – what did we learn?
Outlook – where should further research focus attention?

Please…

3

- Tell me to slow down
- Ask questions as they come to mind
- Ask for clarification if necessary
- Argue my logic
- Eat the food

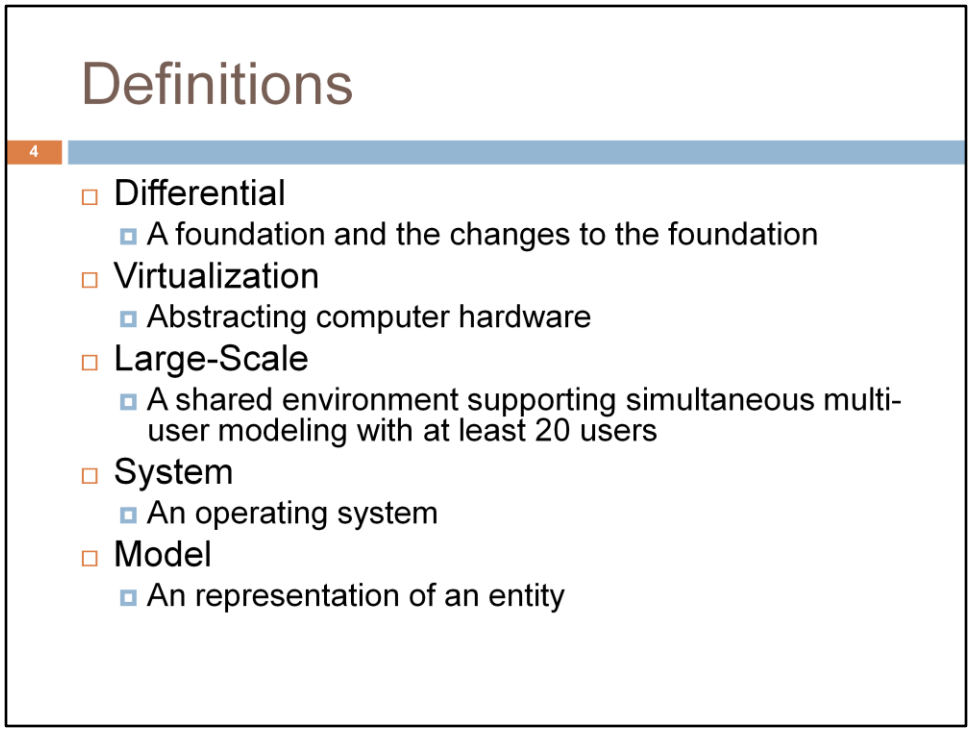

•Virtualization
- •Enables a user to run simultaneously multiple operating systems upon a single computer
- •In this sense its VMware Workstation

•Differential
- •Only save the differences.
- •Analogous to a database– for example, a foundation is defined (as schema) and only differences (as records) are saved RATHER than storing many copies of the database schema, conceptually we store one and store its entries
- •In this sense, its linked clones based on a template

•Model
- •Remember the diorama of the volcano and dinosaurs in 2nd grade? That was a model. It wasn't real. It was a representation of a situation.
- •Remember when you connected the first cable in Netlab or Syslab? Those labs are models because use the labs as a foundation for scientific observation and experimentation.
- •In the sense of this thesis, a system model is an operating system instance, eg virtual machine or an OS installed on a workstation (perhaps via Ghost imaging)

•Large-scale

•Applied systems learning
- •The foundation of mine and many of your bachelors of science in ANSA. This is an entire bachelors program focused around actually playing with technology in order to learn practicality and complexity of technology.

•Labbie
- •started as a labbie in my first quarter at RIT.
- •I suffered through the painstaking manual pre-lab process. A pre-lab is the process by which student workers PREPARED environments prior to student interaction.
- •At one point there were entire teams of students who focused on making these images – across a wide variety of unique hardware combinations. So, this process also included making images every time new hardware comes.
- •When I started, floppy disks contained a ghost.exe and the driver specific to that workstation
- •Then, we changed from room specific floppies to room specific CDs.
- •From there, we went to PXE booting with room specific drivers and universal driver sets.
- •In the three years 2004-2007, there were three different imaging solutions. We kept trying to make student oriented imaging faster!
- •In 2006, RIT bought a new file server to increase throughput!
- •Server Administration - Fully understanding the capabilities of Windows AAA

•580 students – this is one example where we justify having a modeling environment. It is clear that this is important

•Modeling is important to systems administration

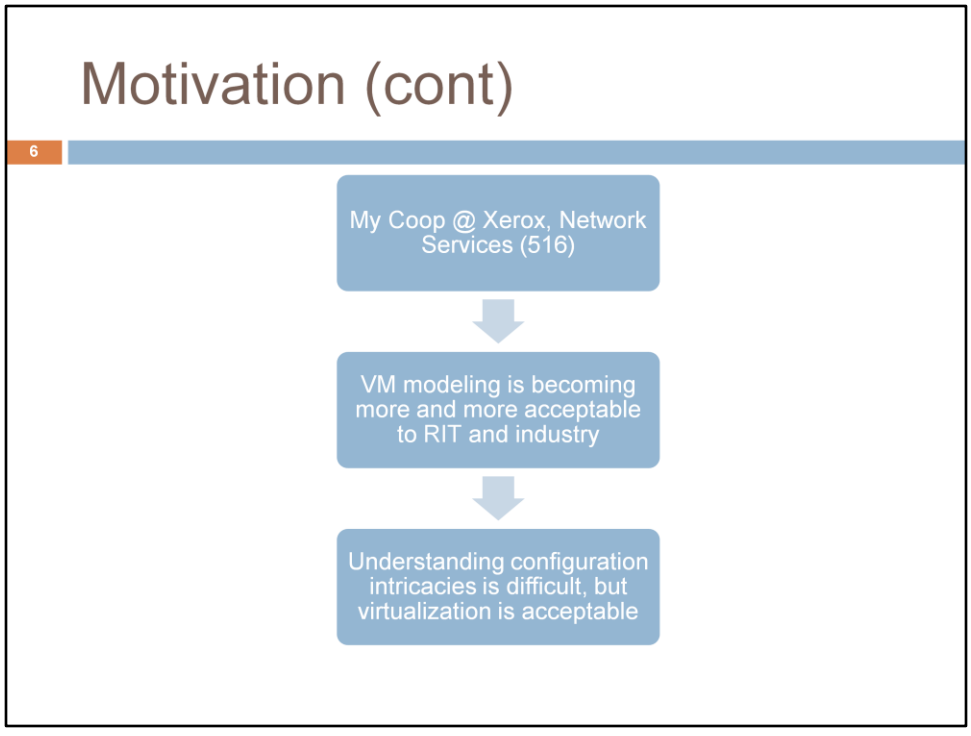

•Coop – Began using VMware in nearly all modeling occurrences
•VM is becoming widely acceptable as a modeling solution for academic and industry – for example, in SANS and Blackhat trainings give trainees a virtual machine. For example, RLES. We teach on the basis that virtualization is a similar enough environment for systems-related material. For example, 585, 886, 760, 761, etc etc.

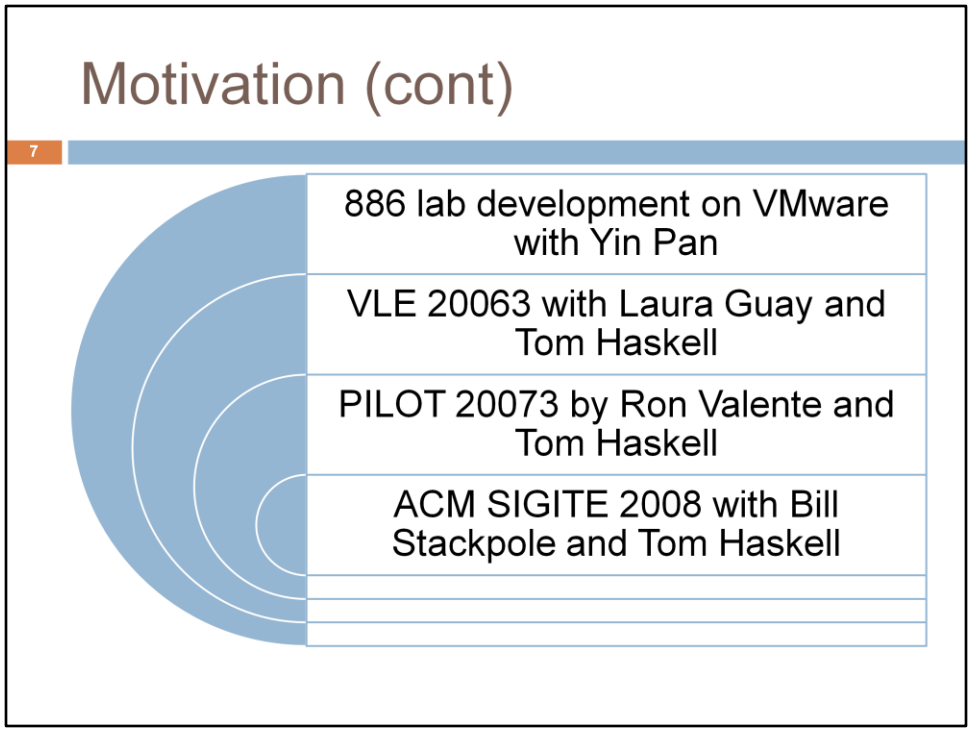

•Worked on virtualization projects
•Laura, Tom and I pushed to have the an environment implemented in 20063
•Ron and Tom revisited the environment by showing faculty a working implementation in 20073.
•Bill, Tom and I compiled our findings into a paper accepted to the ACM Special Interest Group for Information Technology Education
•It is clear that this thesis is a culmination of my efforts at RIT. From Labbie to Grad student, I've always sought to increase efficiency in the lab environment where we model systems to learn.

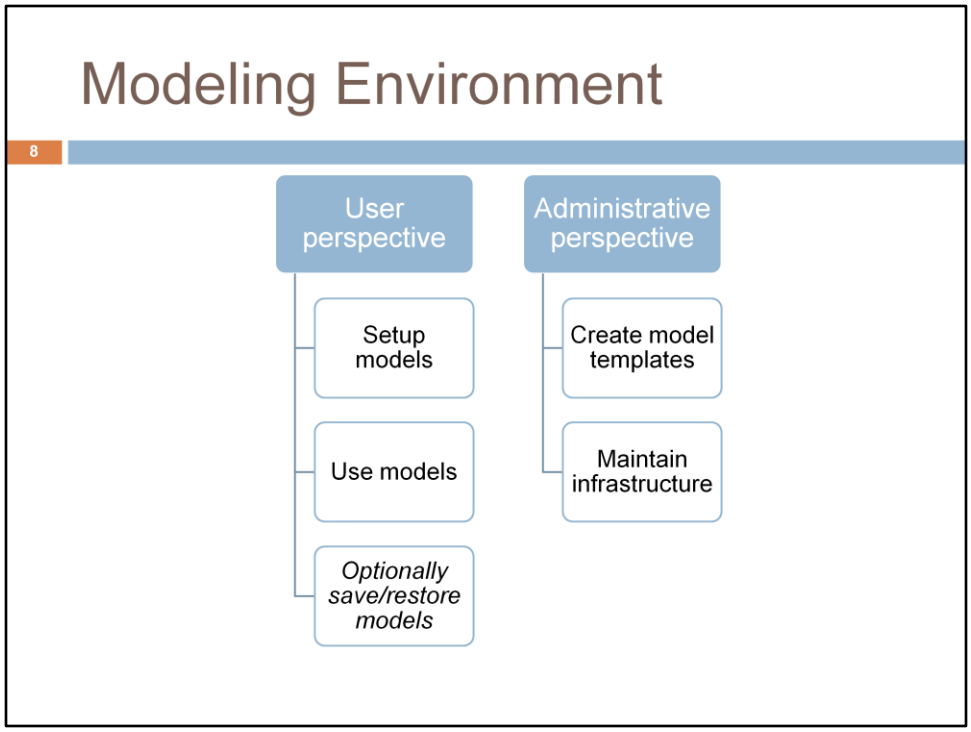

•User Perspective

•Saving/Restoring models is an assumption for this research. It should be easily understood that the ability to save work is beneficial to a process.

Traditional Solutions

9

- Hardware
  - Single workstation

    - Minimal, semi-modeling occurs
    - User can experiment with applications
    - User can observe results
    - In the end, model effects persist after modeling completed

•Single workstation – application/configuration modeling, but modifications persist; therefore we feel the effects of a model after the modeling is done with. For example, I might install IE8 beta to see if it works with my new website, but then I have to deal with having IE8 on my machine!

Traditional Solutions

10

- Hardware

  - Single workstation
    - Minimal, semi-modeling occurs
    - User can experiment with applications
    - User can observe results
    - In the end, model effects persist after modeling completed
  - Multi-boot, similar with a few more OS choices

Multiboot, commonly referred to as dual-boot or triple-boot is a method that enables an OS boot menu where you can select which OS to boot
Only one OS boots in the end

Traditional Solutions

11

- Hardware
  - Single workstation
  - Multi-boot
  - Multiple workstations
    - Install OS on “spare”
    - Experiment and observe effects on “spare”
    - Forget “spare”

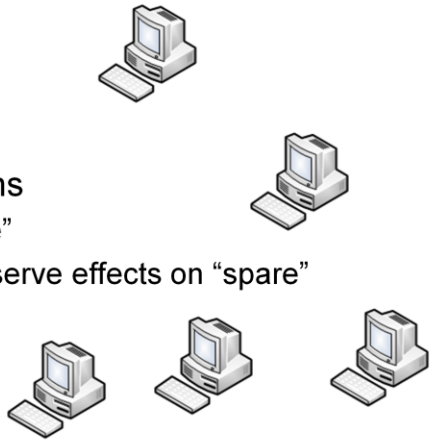

•Extra hardware –

- •full system modeling where modifications are easily erasable.
- •If you have a spare machine, you can just setup the OS twice and configure systems similarly to run tests on a workstation that you “don’t care about”
- •There is still the issue of having to install the OS on the spare
- •Further, saving observable models is difficult unless you have additional hardware

- •We will see that this type of solution doesn’t scale as well as others.

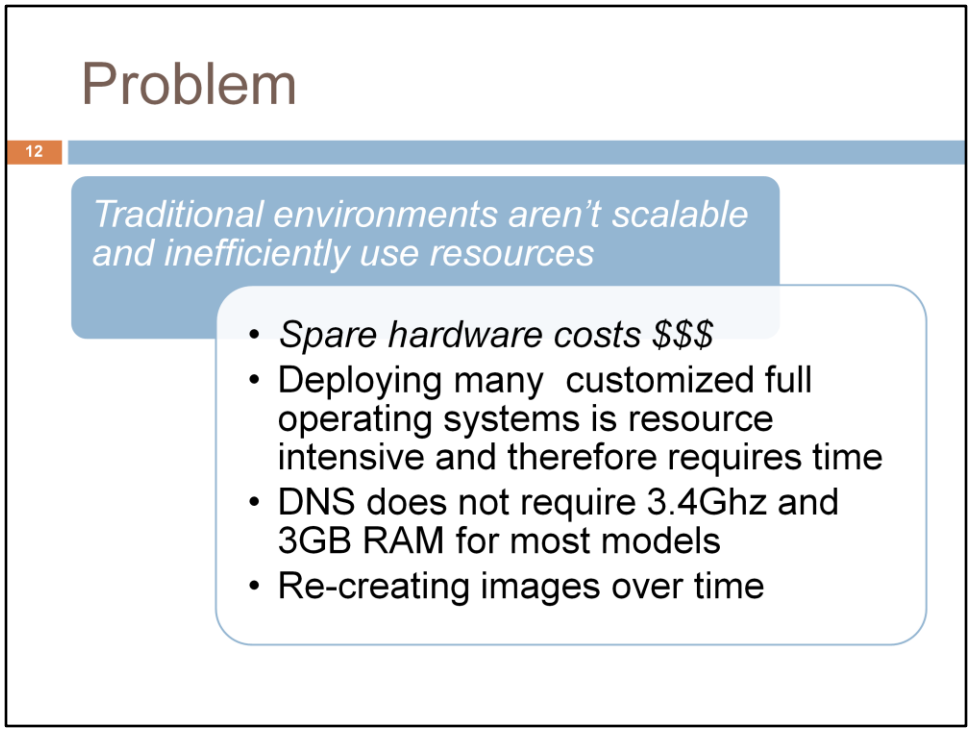

•$$$ - Not all modeling labs are as fortunate as NSSA RIT.

Panacea? Almost…

13

- Virtualization
  - Increases resource utilization by operating multiple operating systems per workstation
  - Decreases dependency on underlying hardware which increases modularity
  - Templates created once and used with many different configurations – basis for differential virtualization
  - **How does one accomplish virtualization?**
    - **Server-driven versus Desktop-driven**

•Server v Desktop

•Still have this notion to decide upon

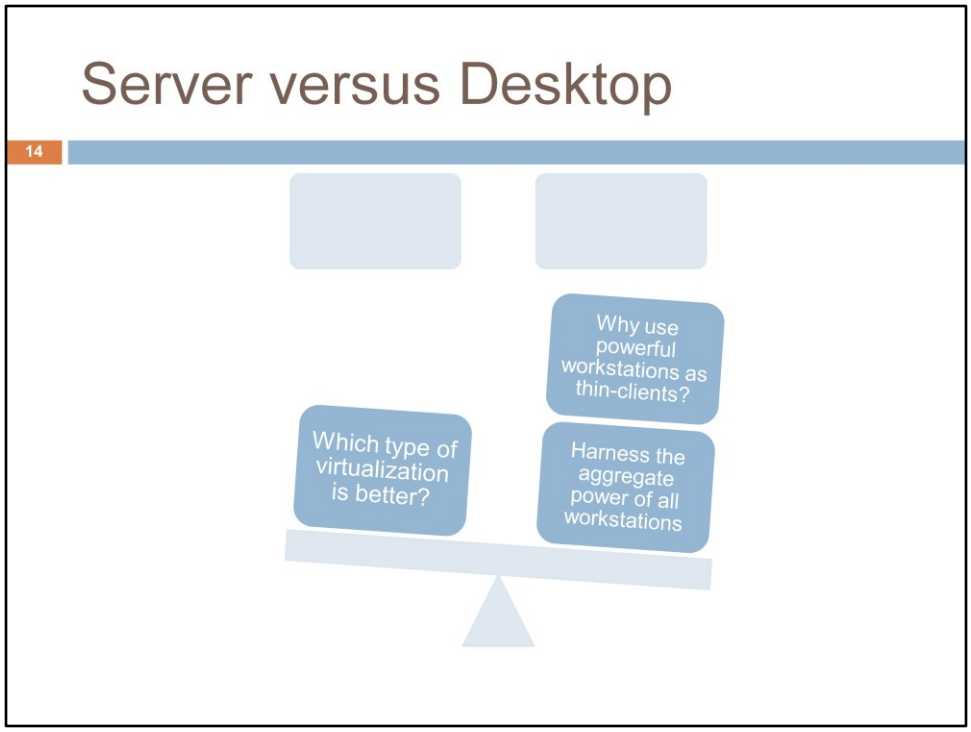

•This research assumes that an many, powerful workstations exist in an organization that requires system modeling
•SYSLAB: 272 GHz CPU with 240 GB memory

Research

15

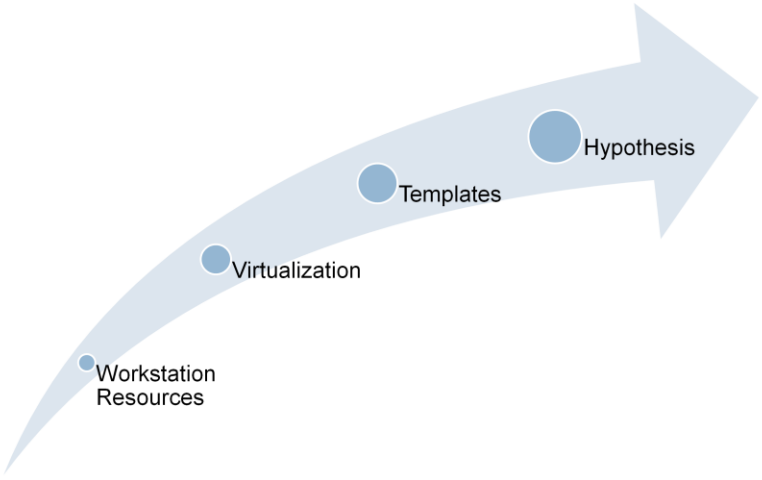
Hypothesis
Templates
Virtualization
Workstation
Resources

Other virtualization solutions

16

- Minimal guest operating system support (solely Linux-based with User-Mode Linux or Microsoft-based with Microsoft Virtual PC)
- Minimal usage of differentiating disks or linked clones
- No physical host image deployment solution
- Claims that VMware costs too much

A new modeling environment

17

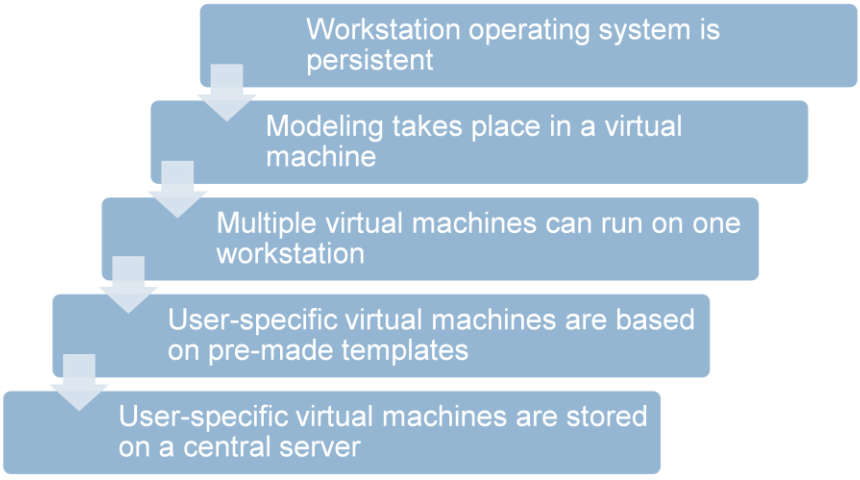

Sample Environment (SYSLAB)

18

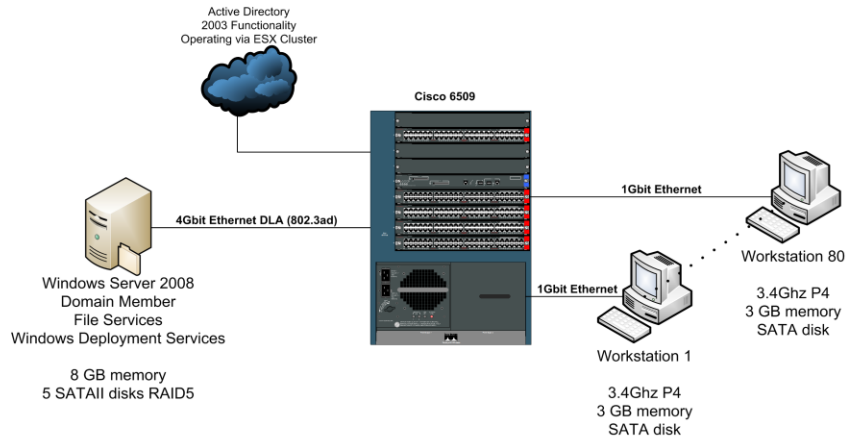

Some assumptions

19

- Preserving a model is beneficial
- Active Directory in-place and used
- More aggregate computing resources in workstations than servers
- VMware Workstation is cheap
  - Free for academia
  - $65 per workstation if at least 250 workstations

A new modeling environment

- Workstation operating system is persistent
- System modeling takes place in a virtual machine
- Workstations operate multiple virtual machines
- Multiple users share multiple workstations
- User-specific virtual machines are based on pre-made templates
- User-specific virtual machines are stored on a central server
  - Storage differs from execution point

A new modeling environment

21

- Specifics:
  - Active Directory domain supports authentication, access control, network services
  - Workstations operate XP or Vista
  - VMware Workstation is installed on workstations
  - VM templates are on workstations
    - E.g. ~90GB templates in SYSLAB
  - Users create linked clones and store them accessible storage solution
    - Storage != Location for execution

VMware was chosen because of the wide guest-operating system support

Hypothesis

22

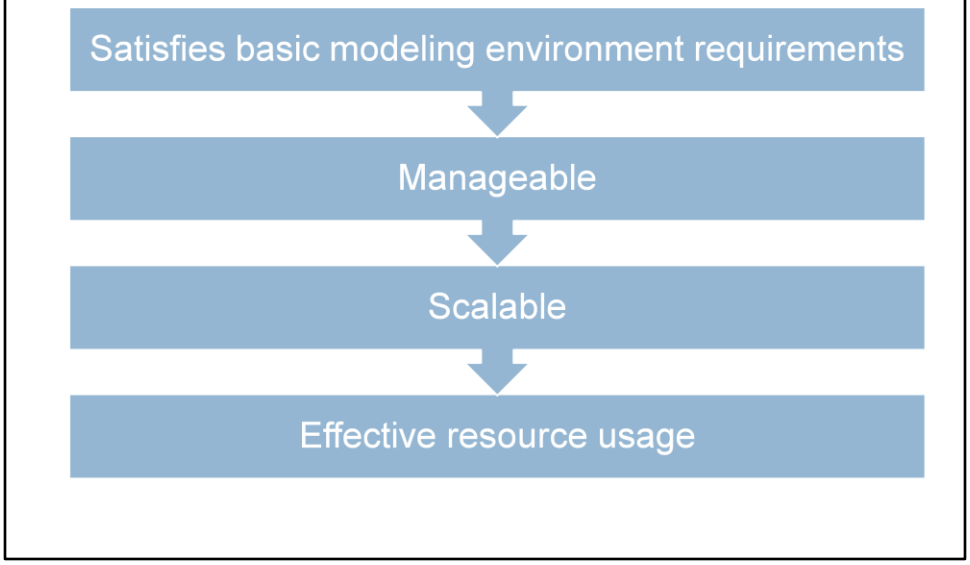

Hypothesis

23

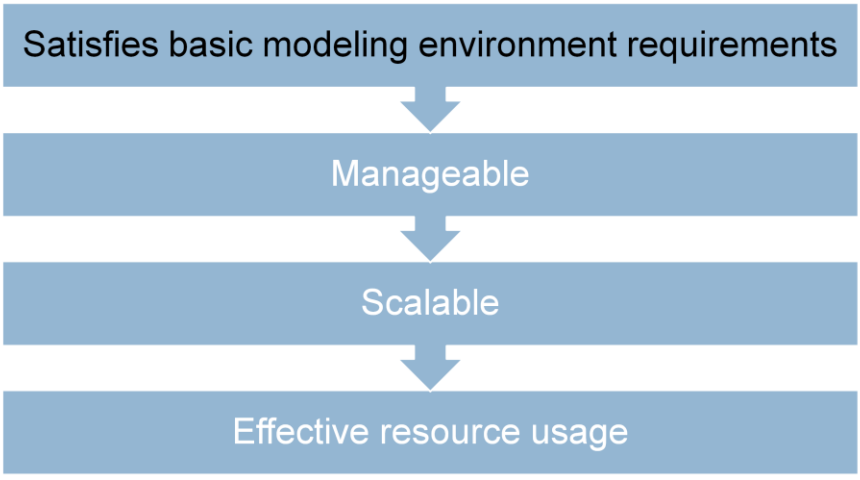

Experiment: Basic Modeling Requirements

- VMware Workstation enables users to…
  - ☑ Setup, use, observe, pause, and record models
  - ☑ Use templates

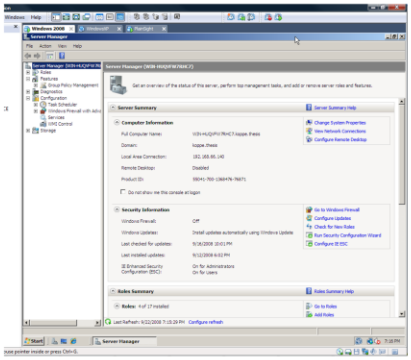

Hypothesis

25

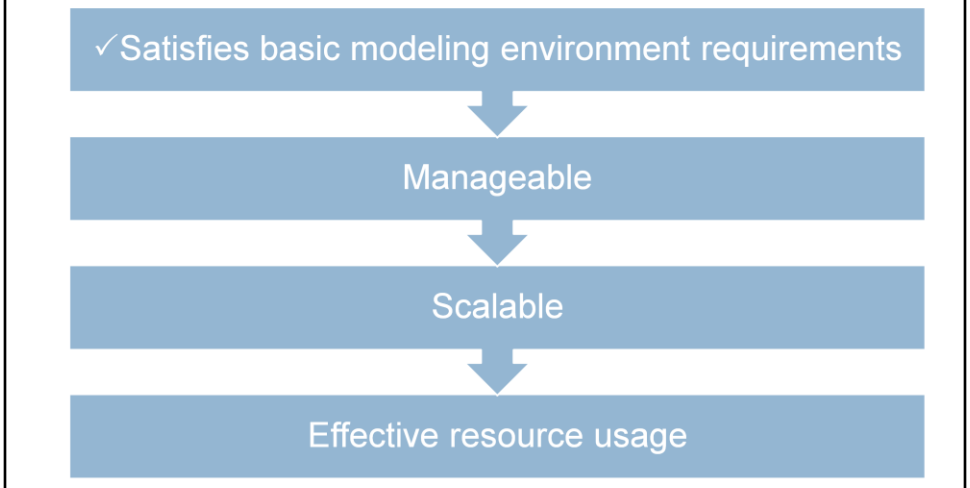

Hypothesis

26

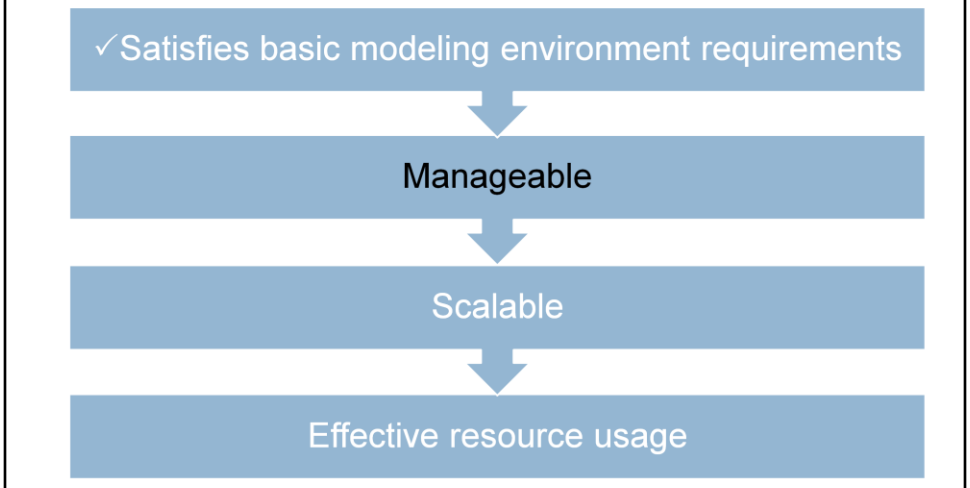

Experiment: Management Tasks

27

- Infrastructure Setup and Maintenance
- Workstation Deployment
- Template Creation
- Template Refresh

Infrastructure Setup and Maintenance

- Microsoft Server facilitates the process
  - Authentication
  - Fine-grained, role-based access control
  - Automation through command line
    - Useradd.pl in document appendix
  - Network services
    - DHCP, DNS roles within MS Server

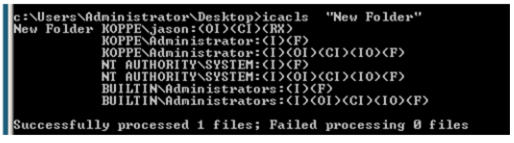

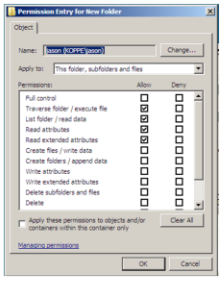

Experiment: Management Tasks

29

- ☑ Infrastructure Setup and Maintenance
- ☐ Workstation Deployment
- ☐ Template Creation
- ☐ Template Refresh

Workstation Deployment

30

- Imaging of some sort is necessary if workstations require operating systems!
- Solution: Windows Deployment Services
  - Install Windows remotely using network booting
  - Multicast
  - Auto-name, auto-join domain
  - Controllable from command line
    - massImage.ps1 in document appendix
  - See document for setup/configuration

How do we get the operating system to each workstation?
Third-party driver accountability
Time-consuming to learn processes – but not much different than learning Ghost or TrueImage processes

Workstation Deployment

31

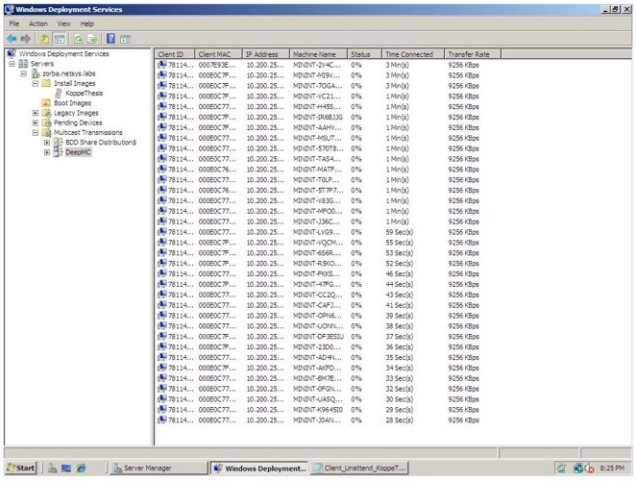

Experiment: Management Tasks

32

- ☑ Infrastructure Setup and Maintenance
- ☑ Workstation Deployment
- ☐ Template Creation
- ☐ Template Refresh

Template Creation

33

- Configure and install an OS to virtual machine
- Create snapshot
- Add one line to .vmx or check one checkbox!

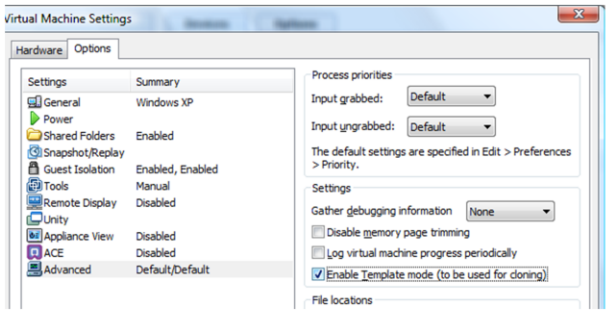

Time consuming process to initialize the repository of templates, but it is a one-time process – no more refreshes when hardware occurs.

Experiment: Management Tasks

34

- ☑ Infrastructure Setup and Maintenance
- ☑ Workstation Deployment
- ☑ Template Creation
- ☐ Template Refresh

Template Refresh

35

- Maintain a template repository
- Uphold consistent templates across all workstations
- Solutions
  - Re-deploy entire OS image new templates
  - Copy templates from server to workstations
    - Differential copy using Robocopy…
    - Multicast using wdsmcast…
  - Differential multicast with Robocopy+wdsmcast
    - Details in document

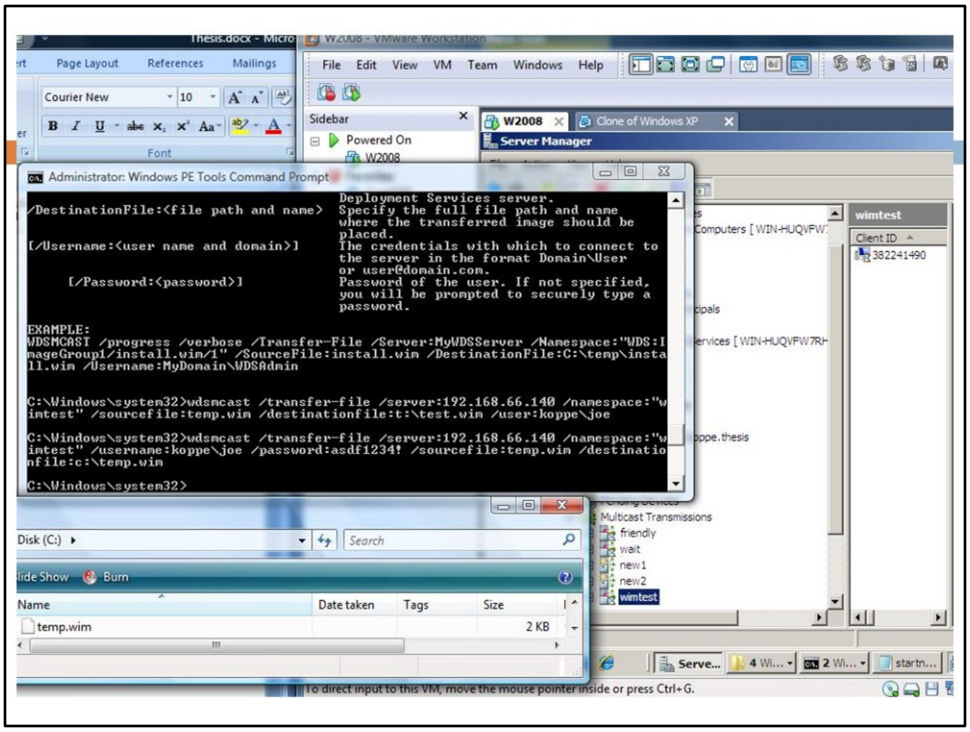
Thesis.docx - Micro
Page Layout
References
Mailings
Courier New
10
Font
W2008 - VMware Workstation
File
Edit
View
VM
Team
Windows
Help
Sidebar
Powered On
W2008
W2008
Clone of Windows XP
Server Manager
Administrator: Windows PE Tools Command Prompt
Deployment Services server.
/DestinationFile:<file path and name>
Specify the full file path and name where the transferred image should be placed.
[/Username:<user name and domain>]
The credentials with which to connect to the server in the format Domain\User or user@domain.com.
[/Password:<password>]
Password of the user. If not specified, you will be prompted to securely type a password.
EXAMPLE:
WDSMCAST /progress /verbose /Transfer-File /Server:MyWDSServer /Namespace:"WDS:ImageGroup1/install.wim/1" /SourceFile:install.wim /DestinationFile:C:\temp\install.wim /Username:MyDomain\WDSAdmin
C:\Windows\system32>wdsmcast /transfer-file /server:192.168.66.140 /namespace:"wimtest" /sourcefile:temp.wim /destinationfile:t:\test.wim /user:koppe\joe
C:\Windows\system32>wdsmcast /transfer-file /server:192.168.66.140 /namespace:"wimtest" /username:koppe\joe /password:asdf1234! /sourcefile:temp.wim /destinationfile:c:\temp.wim
C:\Windows\system32>
wimtest
Client ID
382241490
Multicast Transmissions
friendly
wait
new1
new2
wimtest
Disk (C:)
Search
lide Show
Burn
Name
Date taken
Tags
Size
temp.wim
2 KB
To direct input to this VM, move the mouse pointer inside or press Ctrl+G.

Experiment: Management Tasks

37

- ☑ Infrastructure Setup and Maintenance
- ☑ Workstation Deployment
- ☑ Template Creation
- ☑ Template Refresh

Experiment: Management Tasks

38

- Infrastructure Setup and Maintenance
  - Scriptable user creation/access control
- Workstation Deployment
  - Scriptable deployment configuration
- Template Creation
- Template Refresh
  - Scriptable template modifications

Automation facilitates three of the four tasks

Hypothesis

39

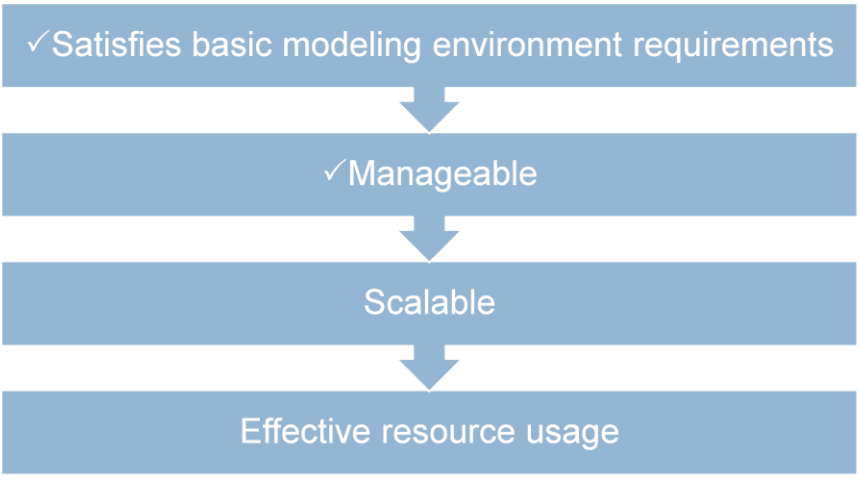

Hypothesis

40

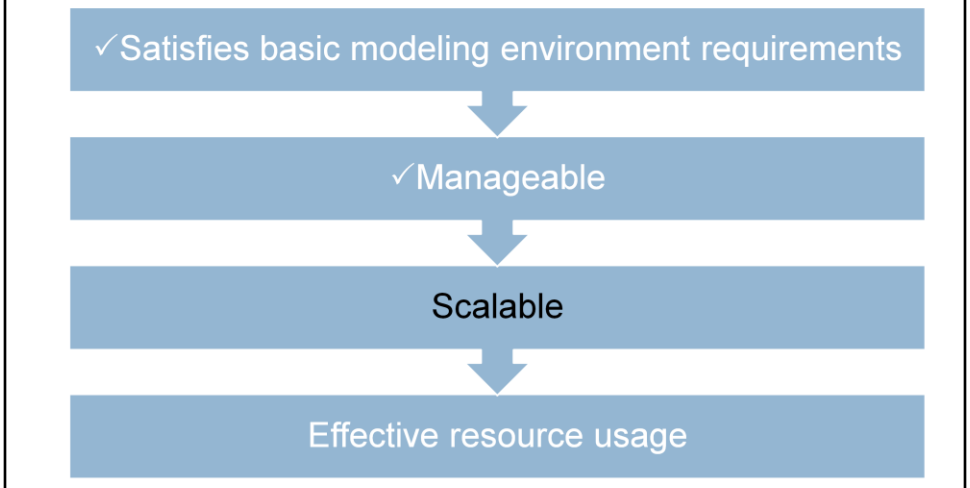

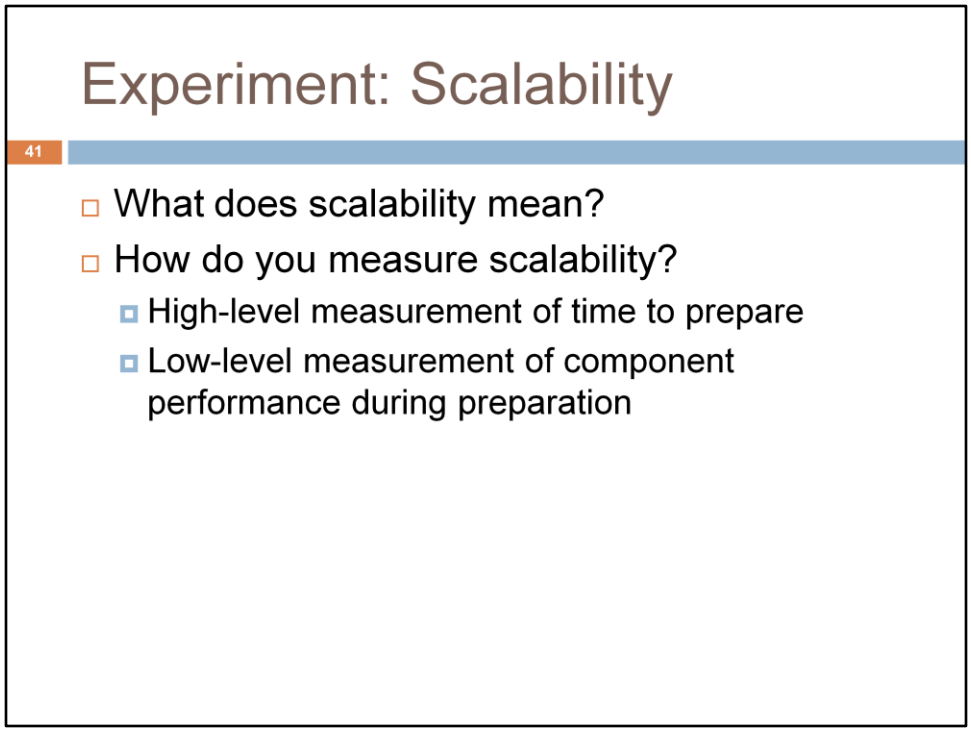

- Scalability
  - as the number of models increases, model setup/preparation/preservation/use is not adversely affected. For example, in the traditional environment, when 80 custom models – ghost images, were being imaged, the whole SysLab suffered.
- Measuring scalability
  - Device performance
  - Process time

Experiment: Scalability

42

- Execution point
  - The location of the linked clone when it is executed by a workstation affects overall system performance
  - Two approaches are evaluated
    - Direct modification: execute while stored on the server
    - Cache-and-update: copy to workstation, execute while stored on workstation, update when done

•Execution point matters – attempted restoration with direct modification of 60 models and failed! Fallback plan is cache-and-update which appears to be less resource intensive.

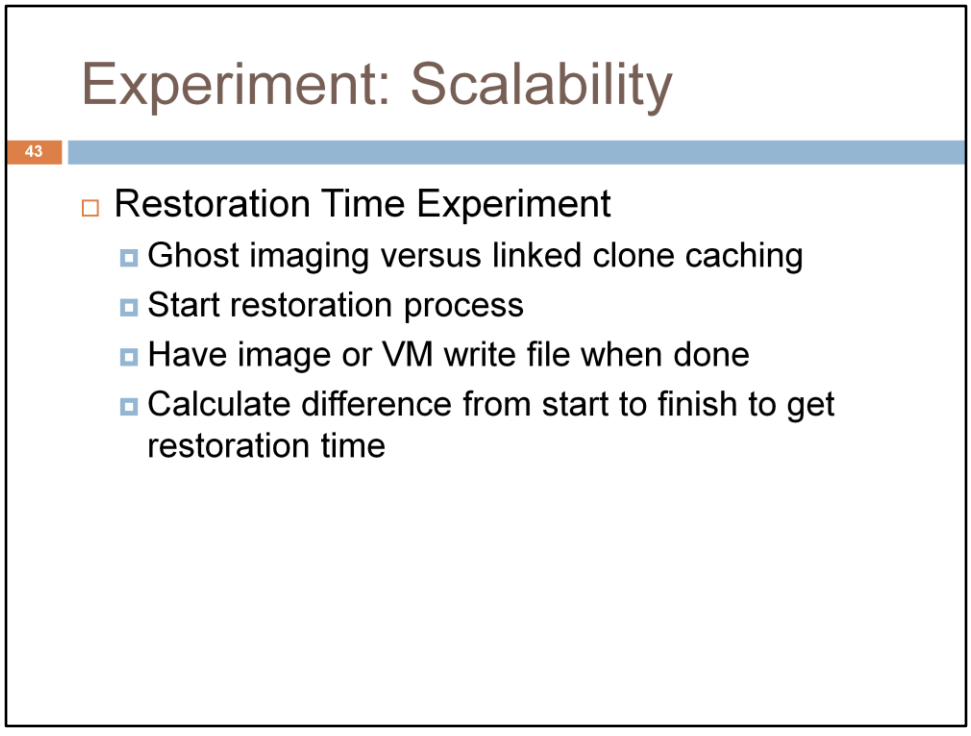

- Scalability
  - as the number of models increases, model setup/preparation/preservation/use is not adversely affected. For example, in the traditional environment, when 80 custom models – ghost images, were being imaged, the whole SysLab suffered.
- Measuring scalability
  - Device performance
  - Process time

Experiment: Scalability

44

- Restoration Time Experiment

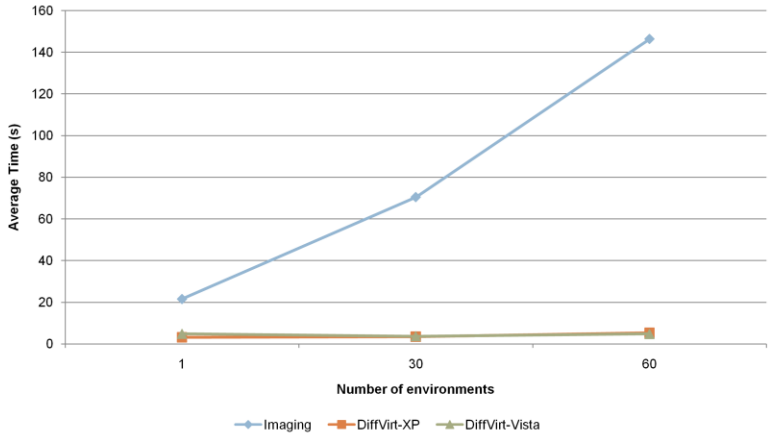

- Scalability
  - as the number of models increases, model setup/preparation/preservation/use is not adversely affected. For example, in the traditional environment, when 80 custom models – ghost images, were being imaged, the whole SysLab suffered.
- Measuring scalability
  - Device performance
  - Process time

Experiment: Scalability

- Restoration Time Experiment (minutes)

	1	30	60
Imaging	21.57	70.51	146.41
DiffVirt-XP	3.17	3.48	5.37
DiffVirt-Vista	4.81	3.68	4.77

Experiment: Scalability

- Direct Modification Characterization
  - Restoration using a direct modification approach with 60 models failed to provide usable models
  - 10 minutes
    - 18/60 restored
    - Of those 18, none were usable

Experiment: Scalability

47

- Direct Modification Characterization
  - Restoration using a direct modification approach with 60 models failed to provide usable models
  - 10 minutes
    - 18/60 restored
    - Of those 18, none were usable
  - Therefore, we should understand what is required for direct modification

Experiment: Scalability

- Direct Modification Characterization
  - Restoration using a direct modification approach with 60 models failed to provide usable models
  - 10 minutes
    - 18/60 restored
    - Of those 18, none were usable
  - Therefore, we should understand what is required for direct modification

Experiment: Scalability

- Direct Modification Characterization
  - 10 minutes: 1.2 billion disk bytes
  - high reads (bottom area), then high writes (top area) – nearly equal overall

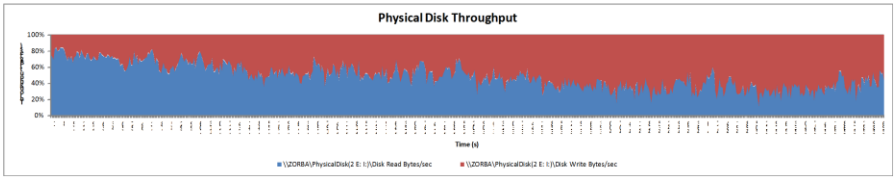

- Caching Characterization
  - 10 minutes: 6 billion disk bytes
  - 99.99% reads

Experiment: Scalability

- Direct Modification Characterization
  - Individual linked clone experiments
    - Create & initially start a linked clone
    - Restore a linked clone
    - Preserve a linked clone (power-off)
  - Measure file server performance
    - Disk throughput
    - Network throughput

Experiment: Scalability

- Direct Modification Characterization
  - Initialize Linked Clone

Workstation OS	Reads	Writes	Read %	Write %
XP	28	4185	*0.7*	***99.3***
Vista	31	1192	*2.6*	***97.4***

Experiment: Scalability

- Direct Modification Characterization
  - Restore Linked Clone

Workstation OS	Reads	Writes	Read %	Write %
XP	3,940	3,804	***50.88***	***49.12***
Vista	7,699	4,598	***62.61***	***37.39***

Experiment: Scalability

- Direct Modification Characterization
  - Shutdown Linked Clone

Workstation OS	Reads	Writes	Read %	Write %
XP	3,037	6,411	*32.14*	***67.86***
Vista	3,664	6,831	*34.91*	***65.09***

Experiment: Scalability

- Direct Modification
  - Overall throughput is greatly reduced in sample environment and renders unusable models
  - Results likely sample-specific
  - Further experimentation should occur to understand which component is the bottleneck
- Cache-and-update
  - Restores 60 models in 7 minutes

Hypothesis

55

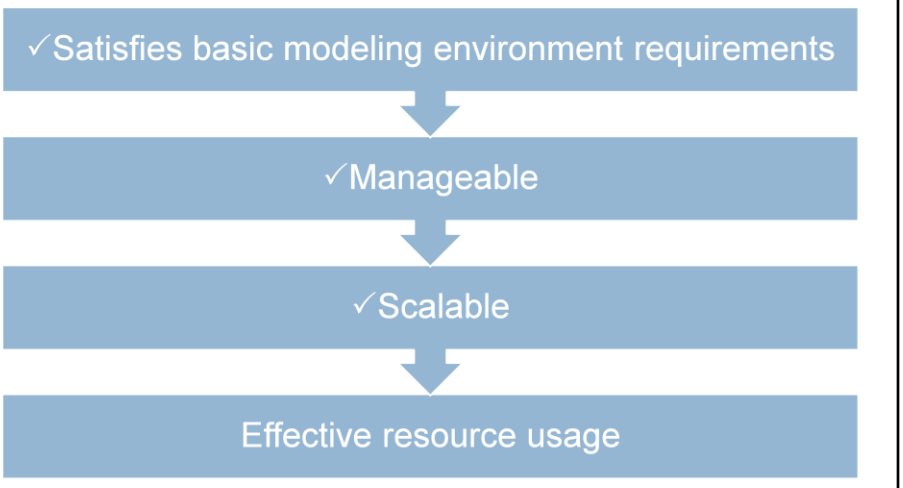

Hypothesis

56

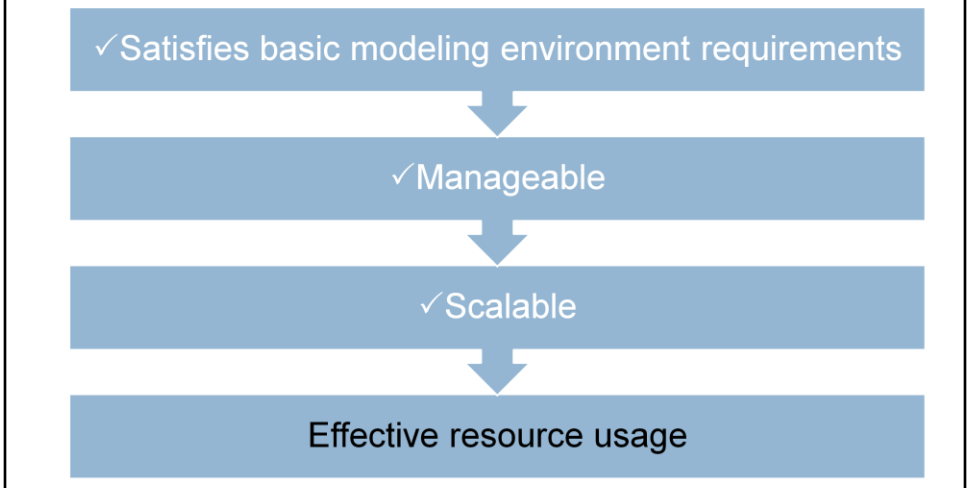

Experiment: Efficient Usage

57

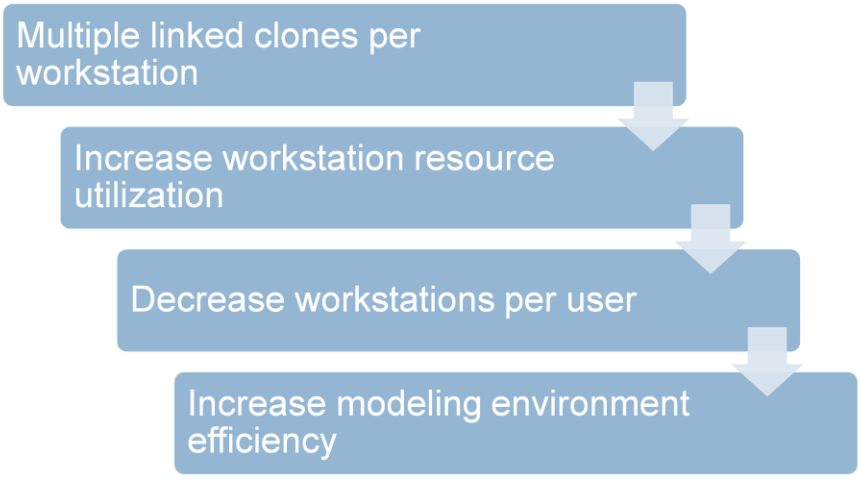

Experiment: Efficient Usage

58

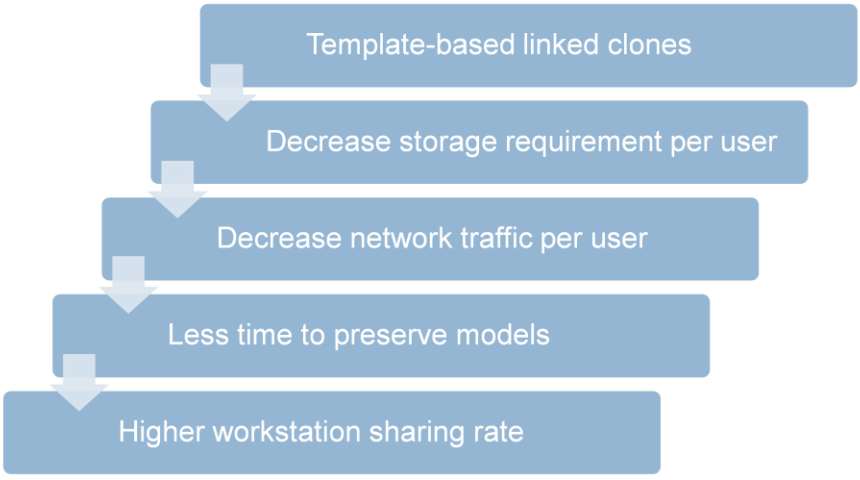

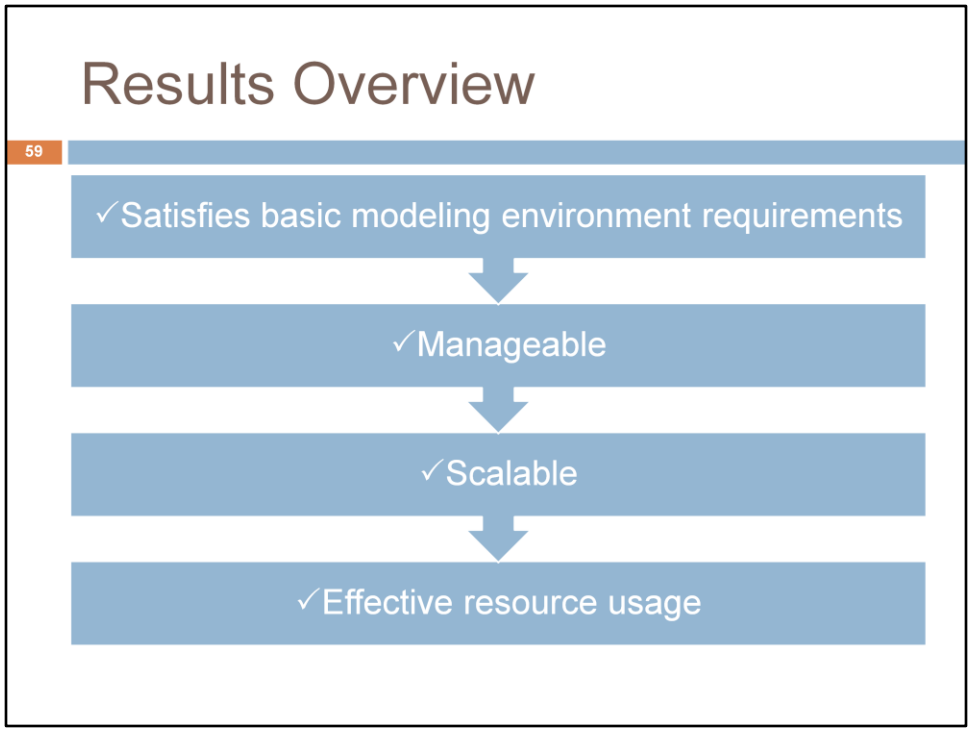

- Satisfies modeling requirements
  - By Using VMware Workstation, users can model entire systems with added benefits
- Manageable
  - Windows is very scriptable!
  - Deploying workstation OS is doable
  - Updating templates is doable
- Scalable
  - Restoration scales with the cache approach, and should work for the update approach as well
- Resource Utilization
  - More lab utilization
  - More resource utilization
  - More resource optimization

Applicability to…

- Development firms
  - Software development/testing
  - OS development/testing
- Research firms
  - Forensics, Malware, Protocols
- Universities
- High Schools?
- Hands-on Certification Workshops

Future Outlook

61

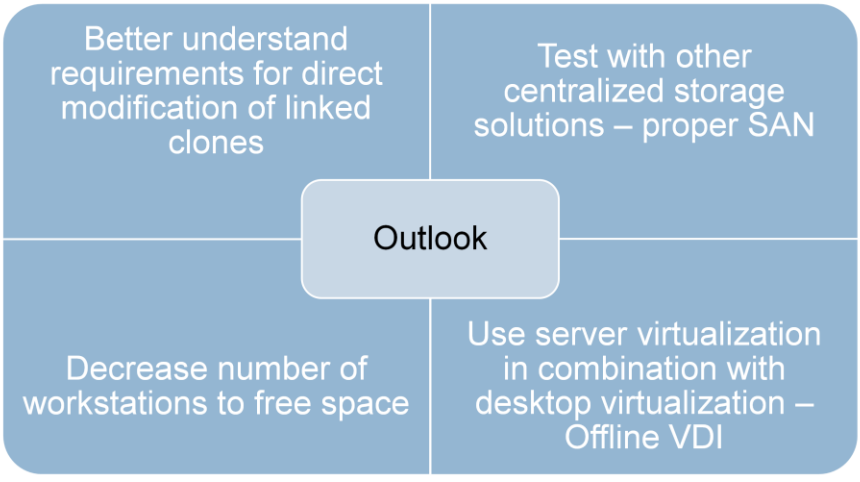

Questions

62

- Would jumbo frames increase throughput?
- When will this environment be replaced and how?
- How do we manage changes to the virtualization platform?
- Would jumbo frames enhance network throughput?
- How could one further measure administrative cost?

Questions and Discussion

63

Jason.Robert.Koppe@gmail.com - http://nssadoc.blogspot.com