A FRAMEWORK FOR CAPACITY ANALYSIS

DEBBIE SHEETZ
PRINCIPAL CONSULTANT
MBI SOLUTIONS

CAPACITY ANALYSIS FRAMEWORK

- What are the essential steps of a Capacity Study?
  1. Obtain the essential question(s) to be answered, the domain, and the time frame for the study
  2. Identify server(s) of interest and their measurements
  3. Analyze historical measurements of the environment
  4. Analyze testing results (if available)
  5. Project future capacity results and/or requirements
- What this isn't about
  - How to do monthly, weekly, etc. capacity reporting
    - Some of what's shown could be used as a basis for regular capacity reporting
  - How to screen a large environment for servers with capacity or performance issues
- Examples show real-world application of the framework (complete capacity report for 3 apps included)
  - Windows and Windows VMs (but methodology is general)

CAPACITY ANALYSIS FRAMEWORK

- **Step 1: Obtain the essential question(s) to be answered, the domain, and the time frame for the study**
  - Identify desired Capacity Thresholds and SLAs
  - Identify source of business forecast
  - Identify capacity people resources to be used in the study
- **Step 2: Identify server(s) of interest and their measurements**
  - Obtain application architecture and application descriptions
  - Identify domain experts
  - Identify data sources (e.g. server measurements, process measurements, business data, etc.)

CAPACITY ANALYSIS FRAMEWORK

- After Steps 1 and 2 have been completed, the answer might be "No, this study can't be done" or
  - "No, this study can't be done in this time frame"
    - This type of study would take x days to complete
  - "No, this study can't be done at all" (due to lack of historical measurements or other required information)
    - Here's a list of the missing measurements
      - Possible approaches to mitigate missing measurements
  - "Yes, there's a higher-level study that can be done in this time frame with the following limitations…"
- Also, negotiation of what the right capacity question to answer may be required at this point

CAPACITY ANALYSIS FRAMEWORK

- **Step 3: Analyze historical measurements of the environment**
  - **Inputs:** Usage, Configuration (cores, memory, processor type, etc.), business volumes, Transaction response times (if available)
  - **Analysis:** Design appropriate workload characterization
  - **Outputs:** Relevant time periods per day/week, relevant business volume periods, cause and effect relationship of business volume and resource usage, most important workload drivers, are performance issues so severe that a capacity analysis can't be performed?

CAPACITY ANALYSIS FRAMEWORK

- **Step 4: Analyze testing results (if available)**
  - **Inputs:** Usage, Configuration, Business volume, Transaction response times (if available)
  - **Outputs:** Compare measured and projected volumes, determine the relationship between simulated load and production loads

CAPACITY ANALYSIS FRAMEWORK

- **Step 5: Project future capacity results and/or requirements**
  - **Inputs:** Identify new hardware and its characteristics
  - **Analysis:** Compare new with existing hardware
  - **Output:** Combine business forecast, capacity thresholds and SLAs, baseline analysis results, hardware characteristics; deliver a presentation and/or report
    - Examples: configuration of VM(s), configuration of physical host(s), number of VMs/hosts required, assignment of VMs to hosts, etc.
      - Server (or VM) configuration
        - Choose the higher of
          - Vendor application requirements (cores, memory, etc.)
          - Usage + projected changes in business volumes, applying desired threshold(s)
      - VM to VMware host ratios
        - VMware designed to dynamically handle over-commitment of resources (CPU and Memory)
        - Capacity planning based on observed and/or projected **usage** (not ratios) assures that adequate physical resources are available
    - Report should have both executive summary and technical content; important assumptions highlighted

STEP 3: ANALYZE HISTORICAL MEASUREMENTS EXAMPLES

- Practical tips
  - When there are multiple types of servers present or a 'what-if' is being evaluated, choose appropriate reporting/modeling units
    - For CPU reporting use a benchmark such as SPECintRate (see CMG 2008 ***Predicting the Relative Performance of CPU*** paper)
      - Avoid using number of cores, CPUs, GHz/MHz, etc.
    - GB/MB for Memory, Disk Space reporting
    - GB/MB per second for disk I/O, network I/O
  - When reporting on VMware VMs, always show the application/OS view of the server (i.e. Windows or Linux) (see CMG 2013 ***Capacity Analysis Techniques Applied to VMware VMs*** paper)
    - VMware/ESX measurements are useful for evaluating the ESX infrastructure

STEP 3: ANALYZE HISTORICAL MEASUREMENTS EXAMPLES

- Practical tips (continued)
  - Resource utilizations are useful only for evaluating past capacity threshold breaches
    - All capacity should be reported combining **configured and used**
  - Select data with granularity matching the stated SLA
    - If SLA is stated for an hour duration, don't use 10 second data!
  - Be sure to understand your measurement data sources and the meaning of the measurements you're using (see CMG 2008 ***Modeling/Sizing Techniques for Different Virtualization Strategies***, and CMG 2010 ***Virtualization Performance and Capacity Data Classification Schema*** papers)

APP A AND B: CAPACITY ANALYSIS

- Migration of applications from Location X to Location Y
  - Loc X: mostly physical (Windows), one virtual server
  - Loc Y: virtual (VMware hosting Windows)
- App B load is a function of
  - Number of transactions which varies by
    - Time of year (business peak)
- Capacity prediction will focus on historical resource utilization (aggregated across all servers)
  - Business cycle is one year
  - Capacity SLA threshold of 70% for CPU and Memory

Statement of utilization-based SLA

All examples headlined in brown text

APP A: CAPACITY DATA

All capacity should be shown as configured vs. used

SPEC benchmark used for all CPU reporting

- CPU Configuration
- 3800 SPEC

- CPU Usage
- 1 year, 230* SPEC

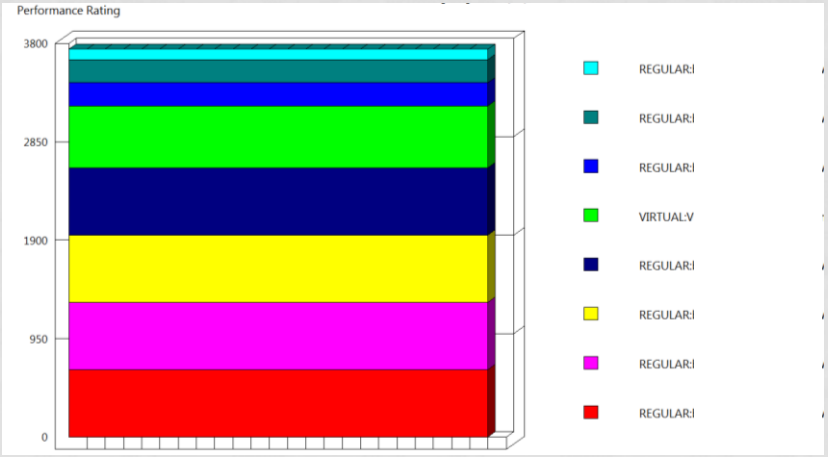

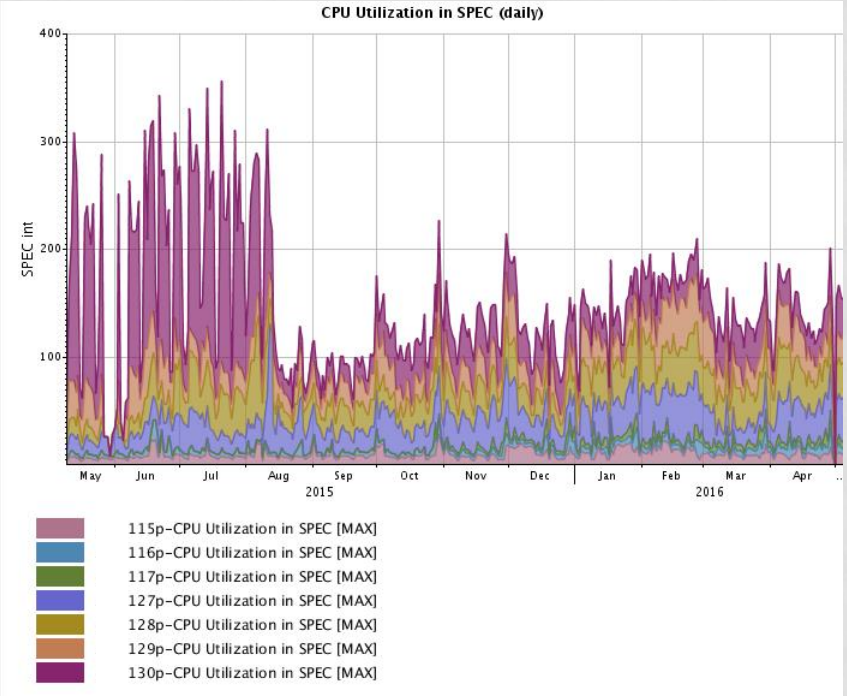

Capacity Risk highlighted

Risk: Usage is not balanced the same as configured capacity

*Ignored May-Aug because code was removed in Aug

APP A: CAPACITY DATA

- Memory Configuration
- 255 GB

GB used for all Memory reporting

- Memory Usage
- 1 year, 81 GB

GB used for all Memory reporting

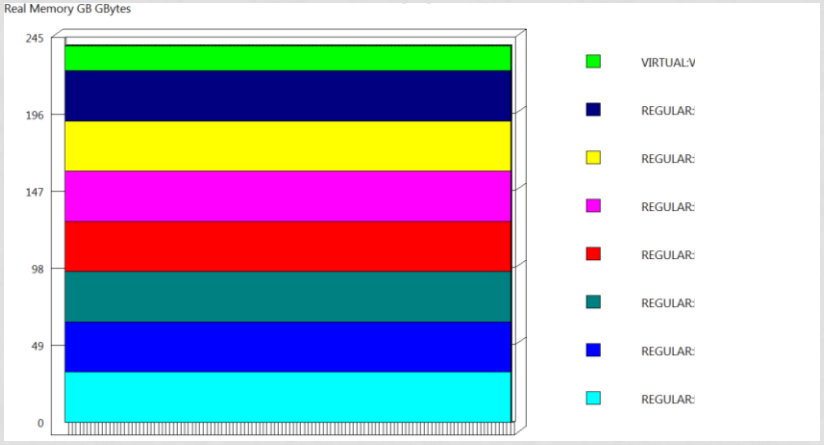

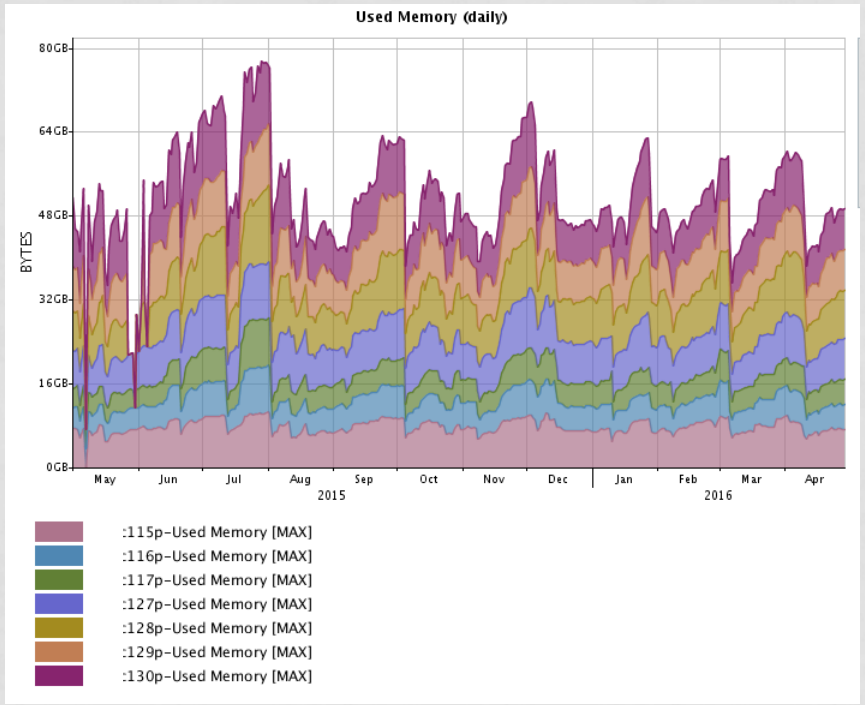

APP A: BUSINESS VOLUME DATA

- Limited (Nov – May) business volume data* (Splunk)

Business volume

Business peak first week of January

*Physical servers only

Usage CPU and Memory

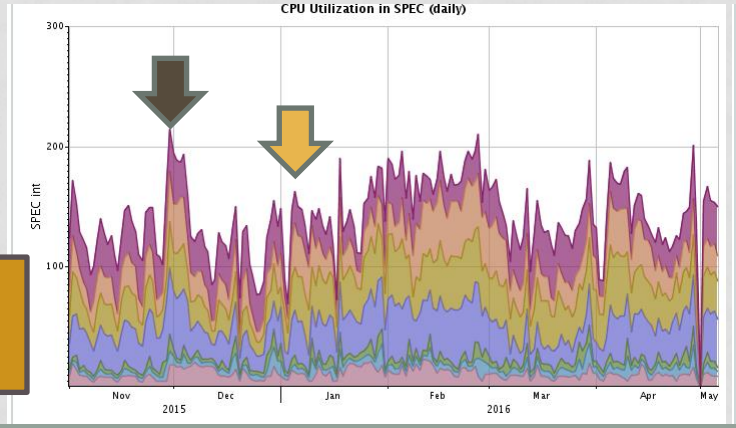

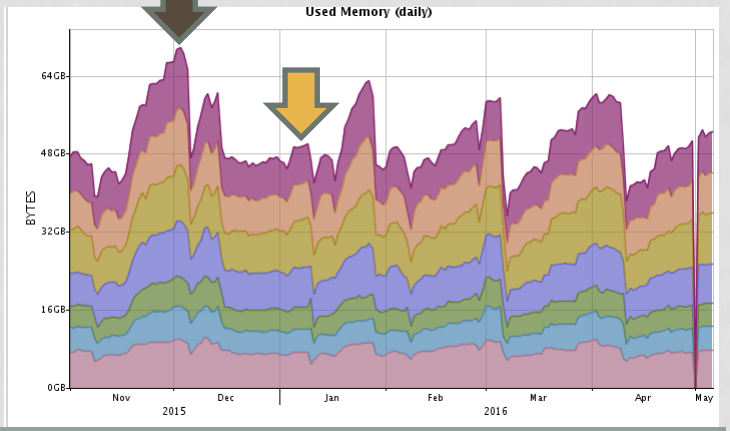

Capacity Risk highlighted

Analysis Risks: No direct correlation between business volume and usage; memory leak behavior is a strong influence on memory usage

APP B: CAPACITY DA

Since the entire application is being moved, aggregated server-level analysis is adequate

- CPU Configuration
- 2900 SPEC

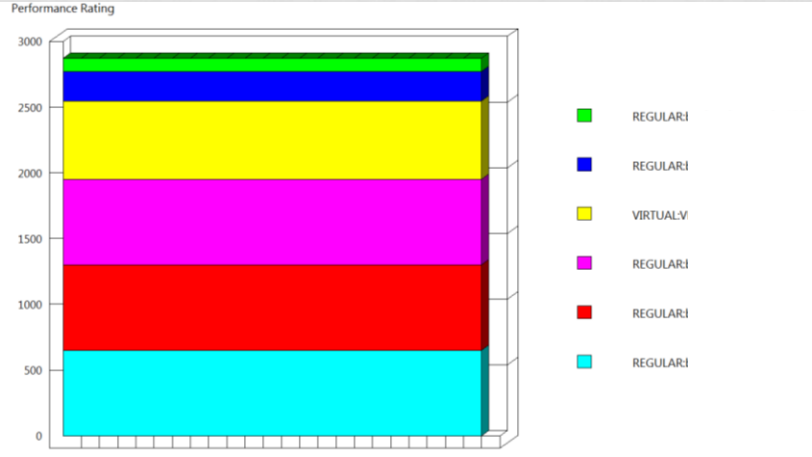

- CPU Usag
- 1 year, 385 SPEC

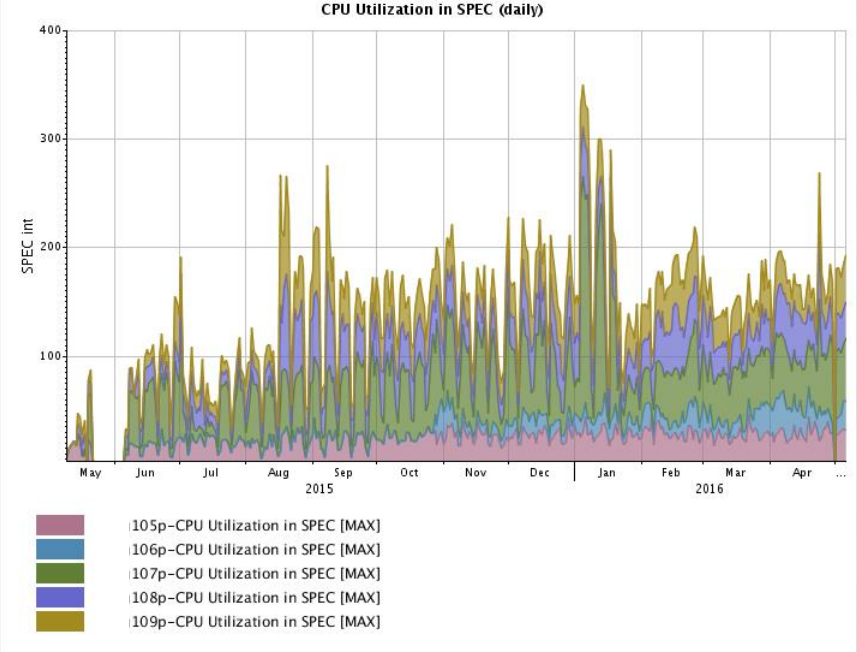

Risk: Usage is not **balanced** the same as configured capacity

APP B: CAPACITY DATA

GB used for all Memory reporting

- Memory Configuration
- 176 GB

- Memory Usage
- 1 year, 122 GB

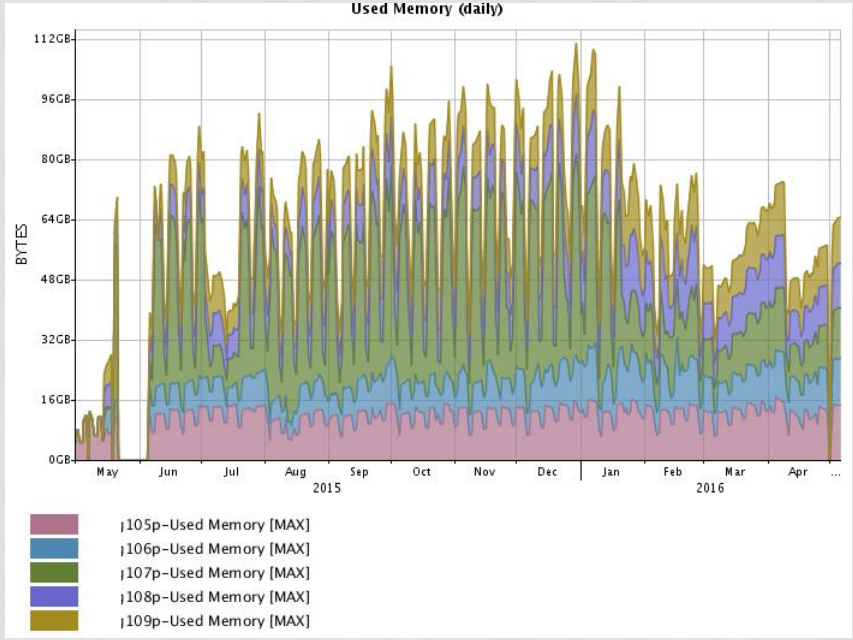

APP B: BUSINESS VOLUME DATA

- Limited (Jan – May) business volume data (Splunk)

Business volume

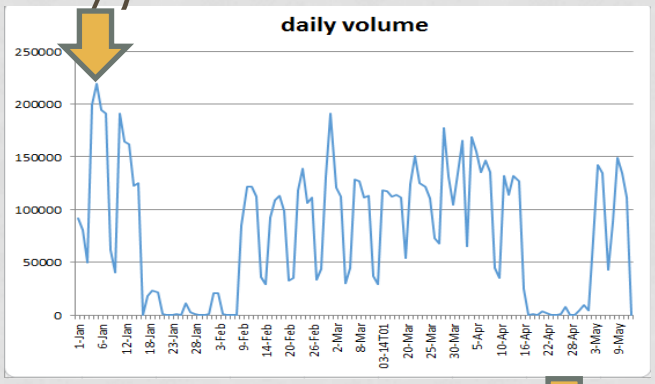

Business peak first week of January

Usage CPU and Memory

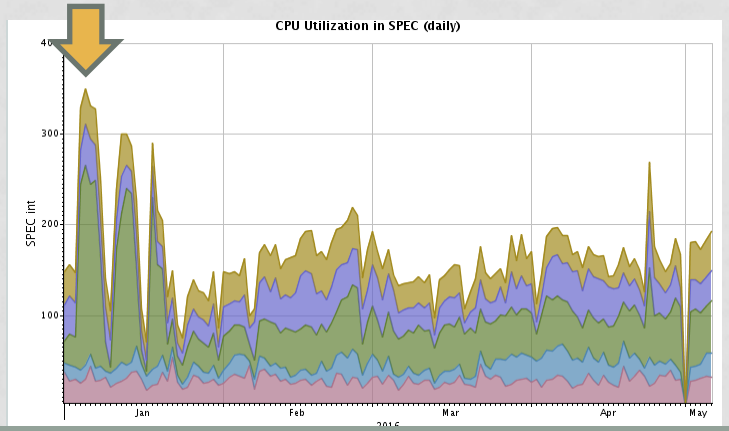

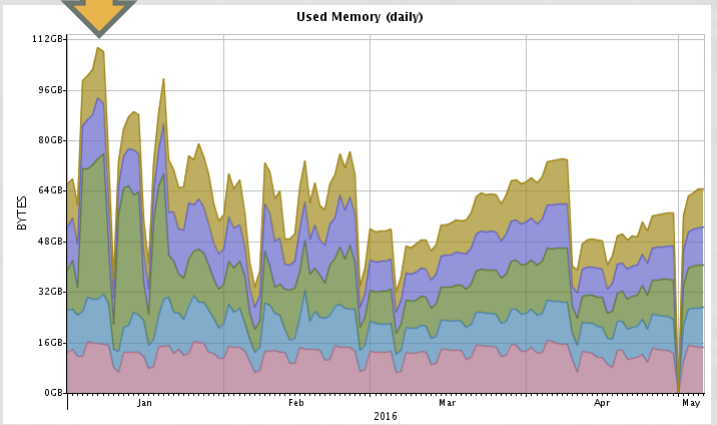

Analysis: Overall correlation between business volume and usage; memory leak behavior is a strong influence on memory usage

APP C: CAPACITY ANALYSIS

- Migration from Location X to Location Y
  - Loc X: mix of physical and virtual (Windows and Vmware)
  - Loc Y: all virtual (VMware hosting Windows)
- App C load is a function of
  - Number of users
  - Work per user, which varies by
    - Time of year (January business peak)
    - Time of day (typical mid-day peak, Monday to Friday)
    - Type of user (4+ types)
- Capacity prediction focuses on number of users (peak), time of day (peak), and time of year (peak)
  - Key element is users per VM
    - Capacity prediction compares projected business volume with projected capacity per VM ➔ number of VMs required to support the peak
  - Capacity SLA threshold of 70% for CPU and Memory

Identification of workload periodicity

Statement of utilization-based SLA

APP C: CAPACITY DATA

- CPU Configuration
- Virtual: 307 SPEC/VM
  - 8 vCPUs per server
  - 38.4 SPEC per vCPU (Intel Xeon E5-2670 @ 2.60GHz)

- CPU Usage
- Oct 2015 – Apr 201
  - Peak: 65% (200 SPEC
    - **Many servers over 70% threshold**
- Feb 2016 - May 2016 is OK
  - Peak: 30% (92 SPEC)

Current CPU capacity is inadequate

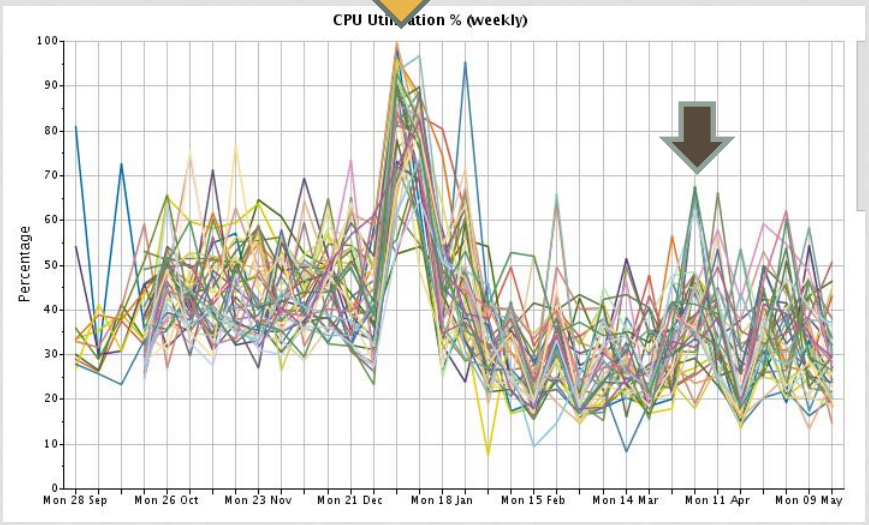

Capacity Risk highlighted

Risk: Usage is not evenly balanced

APP C: CAPACITY DATA

- Memory Configuration
- Virtual: 64 GB /VM

- Memory Usage
- 1 year
  - Utilization 30% – 60%

Current Memory capacity is adequate

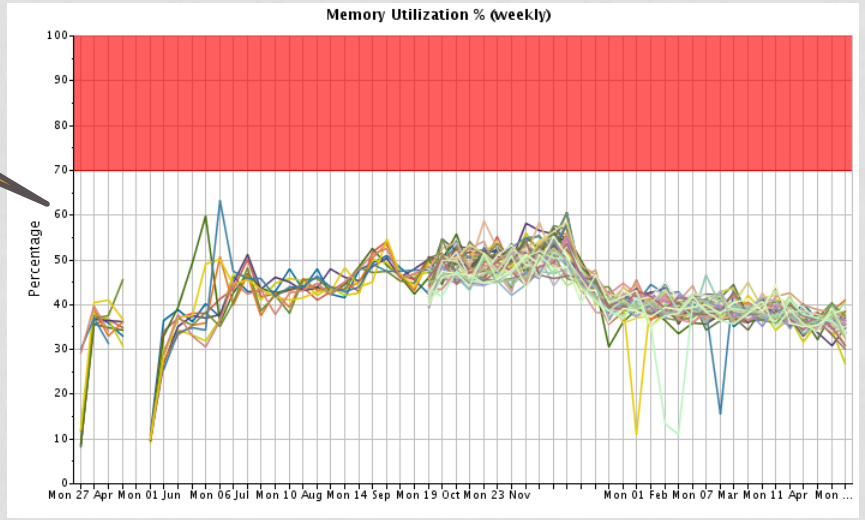

APP C: BUSINESS VOLUME DATA

- Server load is a combination of number of users and what kind of work they are doing (Splunk data)

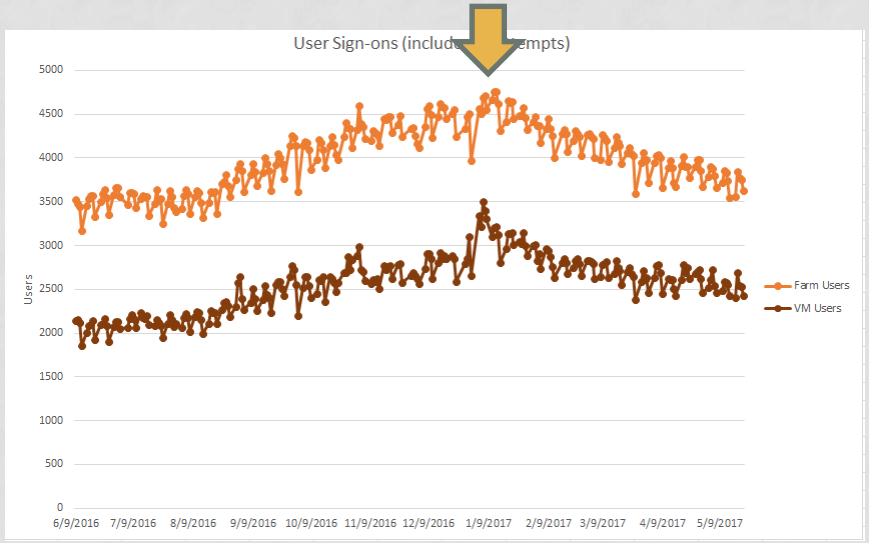

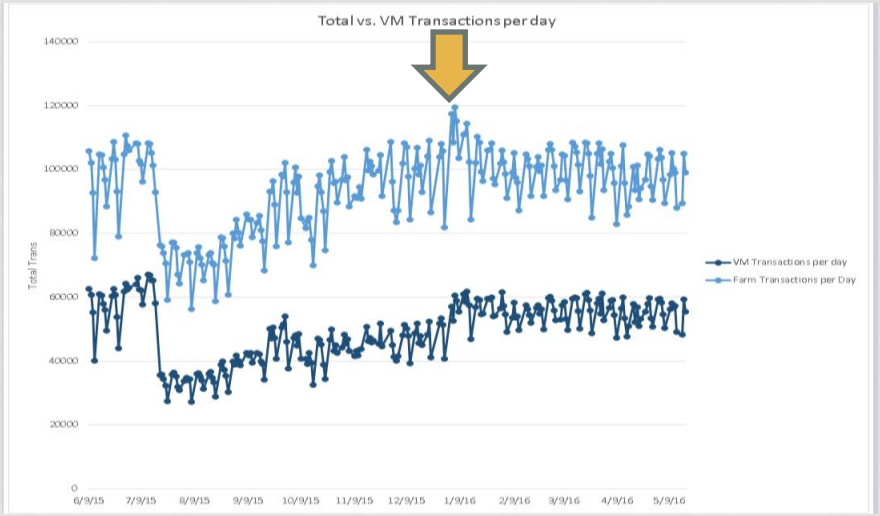

Business volume
All App C vs.
VM App C only

Users
All App C vs.
VM App C only

APP C: BUSINESS VOLUME DATA

- Server load is a combination of number of users and what kind of work they are doing (Splunk data)

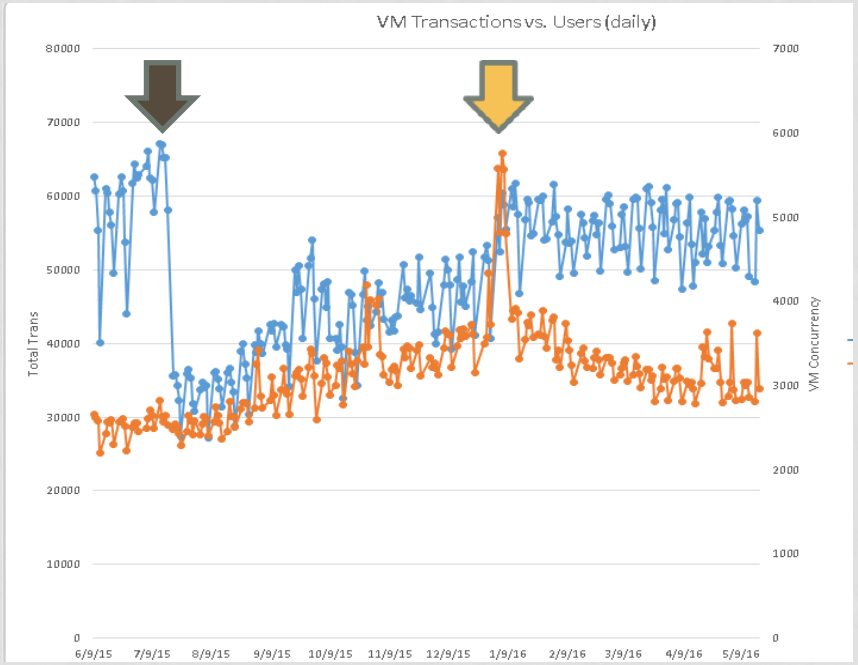

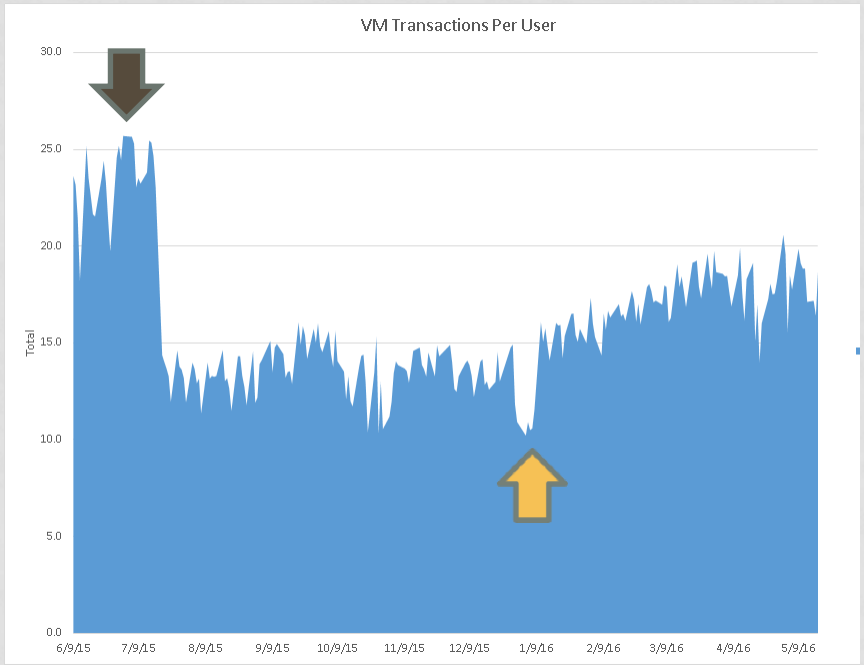

Transactions per user varies depending on the time of year; peak number of users generate fewer transactions per user

APP C: BUSINESS VOLUME DATA

- June 2015 – May 2016 business volume data (Splunk)

Business volume

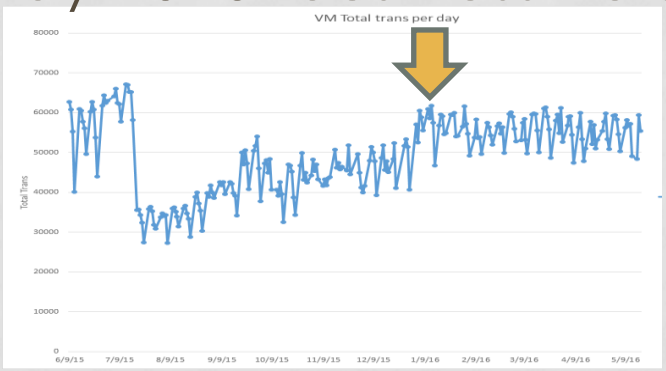

Business peak first week of January

Usage CPU and Memory

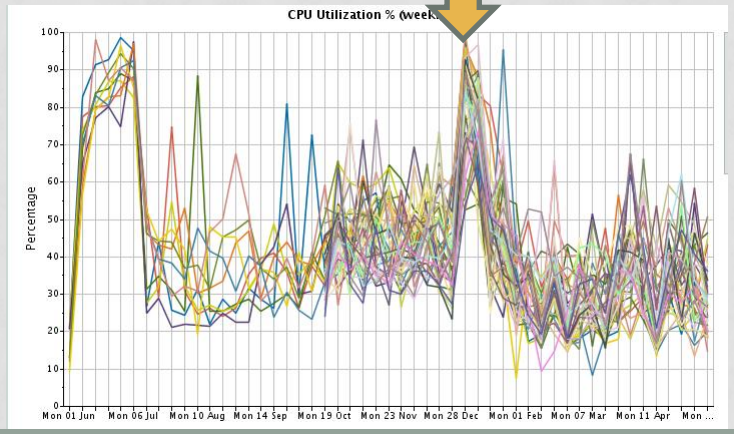

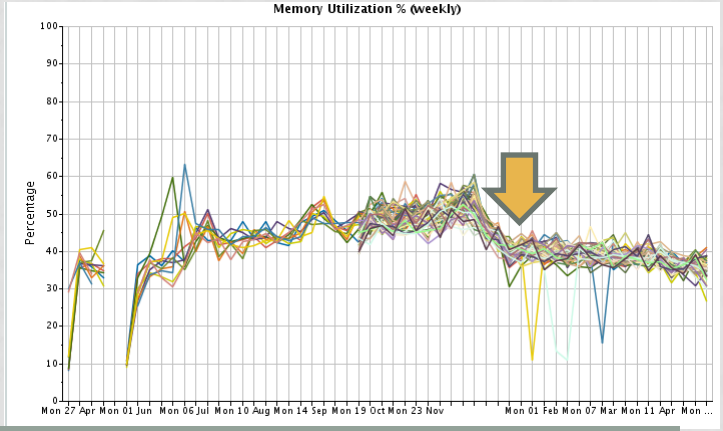

Analysis : Some correlation between business volume and CPU usage; but not memory usage

STEP 4: ANALYZE TESTING RESULTS EXAMPLES

- Practical tips
  - Sometimes testing results aren't needed
  - Most useful when "large" resource changes are being evaluated and it's a production application
  - Biggest challenge is duplicating the production workload successfully
    - Requires detailed understanding of what makes up the production workload
    - Requires ability to duplicate the transactions (at least the most important ones) and their environment in test
    - Requires relevant hardware in the test environment

APP A: TEST RESULTS

- Preliminary results don't mimic production
  - Not all transaction types represented (5 out of 11)
    - Missing types could be influence results substantially
  - No time to get proportions of types correct
- 8 VMs (same as migration configuration)
  - 8 vCPUs and 8 GB memory
- Memory usage is lower than what's been measured in production
  - No conclusion as to what's the cause
    - VMs are being rebooted frequently
    - VMs just use less than physical servers?

Bottom line is that these results can't be used

Analysis Risk: Test results are **inconclusive**

APP B AND C: TEST RESULTS

- None have been made available

If management expected testing, it didn't happen

STEP 5: FUTURE CAPACITY RESULTS EXAMPLES

- Practical tips
  - Trending of historical resource measurements is not a substitute for a business forecast
    - Tracking of resource utilization trends shouldn't be done unless you're sure that the resource configuration **has not changed**
  - When there are multiple types of servers present or a 'what-if' is being evaluated, choose the right units
    - For CPU reporting use a benchmark such as SPECintRate (see CMG 2008 ***Predicting the Relative Performance of CPU*** paper)
      - Avoid using number of cores, CPUs, GHz/MHz, etc.
    - GB/MB for Memory, Disk Space reporting
    - MB/sec for disk I/O, network I/O
  - For VMware (or other virtualization platforms), base predictions on actual/projected usage, not ratios

Should be the first slide; may be the last slide for the executive!

APP A, B, AND C CAPACITY RECOMMENDATION SUMMARY

Important assumptions highlighted

- Summary of capacity recommendations
- Important study requirements have not been met
  - Limited testing results from new environment
  - Business volume data
    - Limited history available
    - No business forecasts

High-level summary of what's recommended

	VM CPU	VM MEMORY	HOST CPU	HOST MEMORY
App A	OK	OK*	OK	OK*
App B	OK	OK*	OK	OK*
App C	INCREASE	INCREASE	INCREASE	INCREASE

*Earlier proposed configuration was too small

APP A: COMPARISON OF CAPACITY

A. Statement of available capacity

- CPU
  - Loc X: 3800 SPEC
    - 8 servers
  - Loc Y: 3072 SPEC
    - 10 servers @ 307 SPEC
      - 8 vCPUs per server
      - 38.4 SPEC per vCPU (Intel Xeon E5-2697 v2 @ 2.70GHz)
  - **Loc Y is 20% smaller**
- MEMORY
  - Loc X: 255 GB
    - 8 servers
  - Loc Y: 160 GB
    - 10 servers
      - 16 GB per server
  - **Loc Y is 40% smaller**

High-level summary of what was found

APP A: RECOMMENDED VM CONFIGURED CAPACITY: APPLICATION

B. Statement of application-required capacity

- CPU
  - App vendor recommendation
    - none
- MEMORY
  - App vendor recommendation
    - 16 (16 active, 16 standby) engines **per server**
      - Loc X: 128 Engines
      - Loc Y: 160 Engines
  - 175 MB usage per engine pair (125 MB + 50 MB)
  - Application memory leak
  - Loc Y: 2.8 GB (app) + 3.5 GB (Windows, etc.) per server
    - Using threshold of 70%, you need 12 GB

Capacity Risk highlighted

Capacity Risk: Horizontal scaling of VMs doesn't mitigate lack of application memory because memory utilization isn't a function of transaction load

APP A: RECOMMENDED VM CONFIGURED CAPACITY: USAGE

C. Statement of usage-based capacity

- CPU
  - No business forecast, so using historical peak
    - Loc X: 230 used of 3800 SPEC ➔ 6% utilization
    - Loc Y: 230 used of 3072 SPEC ➔ 8% utilization
  - So the lesser configuration is adequate from a usage perspective
- MEMORY
  - No business forecast, so using historical peak
    - Loc X: 81 used of 255 GB ➔ 32% utilization
    - Loc Y: 81 used of 160 GB ➔ 50% utilization
  - Loc Y is adequate
    - Better than required to meet threshold of 70%
      - 12 GB per server

APP A: RECOMMENDED CONFIGURED CAPACITY: SUMMARY

D. Conclusion taking into account capacity available, and required

- CPU
  - Proposed 8 vCPU configuration is adequate
  - Since utilization will be very low, it's possible to over-commit for these VMs
- MEMORY
  - Proposed 16 GB configuration is adequate
    - Usage requires 12 GB
    - Application requires 12 GB
  - Since utilization will be low, it's possible to over-commit for these VMs

Capacity Risk highlighted

Capacity Risk: Usage must be **balanced** across VMs to use horizontally-scaled configured capacity; production currently unbalanced; recommend restarting application weekly for memory

APP B: COMPARISON OF CAPACITY

- CPU
  - Loc X: 2900 SPEC
    - 6 servers
  - Loc Y: 2500 SPEC
    - 8 servers
      - 8 vCPUs per server
      - 38.4 SPEC per vCPU (Intel Xeon E5-2697 v2 @ 2.70GHz)
  - **Loc Y is 14% smaller**
- MEMORY
  - Loc X: 176 GB
    - 6 servers
  - Loc Y: 384 GB
    - 8 servers
    - 48 GB per server
  - **Loc Y is 118% larger**

APP B: RECOMMENDED VM CONFIGURED CAPACITY: USAGE

- CPU
  - No business forecast, so using historical peak
    - Loc X: 385 used of 2900 SPEC ➔ 13% utilization
    - Loc Y: 385 used of 2500 SPEC ➔ 15% utilization
  - So the new configuration is adequate from a usage perspective
- MEMORY
  - No business forecast, so using historical peak
    - Loc X: 122 used of 176 GB ➔ 70% utilization
    - Loc Y: 122 used of 384 GB ➔ 32% utilization
  - Loc Y is larger so no issue here
    - 24 GB per server

APP B: RECOMMENDED VM CONFIGURED CAPACITY: APPLICATION

- CPU
  - App vendor recommendation
    - none
- MEMORY
  - App vendor recommendation
    - 24 engines (24/24 A/S) (37/37 for one) **per server**
      - Loc X: 181 Engines
      - Loc Y: 205 Engine
  - 250 MB usage per engine pair (200 MB + 50 MB)
  - Application memory leak
  - Loc Y: 6 GB (app) + 3.5 GB (Windows, etc.)
    - Using threshold of 70%, you need 16 GB

Capacity Risk highlighted

Capacity Risk: Horizontal scaling doesn't mitigate lack of application memory because memory utilization is not a function of transaction load

APP B: RECOMMENDED CONFIGURED CAPACITY: SUMMARY

- CPU
  - Proposed 8 vCPU configuration is adequate
  - Since utilization will be low, it's possible to over-commit for these VMs
- MEMORY
  - Proposed 48 GB configuration is adequate
    - Usage requires 24 GB
    - Application requires 16 GB
  - Since utilization will be low, it's possible to over-commit for these VMs

Capacity Risks highlighted

Capacity Risk: Usage must be **balanced** across VMs to use horizontally-scaled configured capacity; need to restart application weekly for memory

APP C: COMPARISON OF CAPACITY

- CPU
- Loc X: 30011 SPEC
  - Virtual: 19648 SPEC
    - 64 servers @307 SPEC
      - 8 vCPUs @38.4 SPEC per vCPU (Intel Xeon E5-2670 @ 2.60GHz)
  - Physical: 9851 SPEC
    - 11 G1 @101 SPEC
    - 10 G7 @224 SPEC
    - 10 G8 @650 SPEC
- Loc Y: 30240 SPEC
  - 96 servers @307 SPEC
    - 8 vCPUs @38.4 SPEC per vCPU (Intel Xeon E5-2670 @ 2.60GHz
- Loc Z: 7982 SPEC (future spare capacity)
  - 26 servers@307 SPEC
    - 8 vCPUs @38.4 SPEC per vCPU (Intel Xeon E5-2697 v2 @ 2.70GHz
- **Loc Y is 1% larger**
  - **per VM identical**

- MEMORY
- Loc X: 5728 GB
  - Virtual: 4096 GB
    - 64 servers @64 GB
  - Physical: 1632 GB
    - 11 G1 @32 GB
    - 10 G7 @64 GB
    - 10 G8 @64 GB
- Loc Y: 6144 GB
  - 96 servers @64 GB
- Loc Z: 1664 GB (future spare capacity)
  - 26 servers @64 GB
- **Loc Y is 7% larger**
  - **per VM identical**

APP C: RECOMMENDED CONFIGURED CAPACITY: SUMMARY

Summary is shown first

- CPU
  - Continue with 8 vCPU configuration
  - Number of VMs needs to be increased
    - January: 158
    - Near-term: no business forecast to use for sizing
  - Continue work profiling types of transactions and types of users
    - Most likely user siloing will be required
- MEMORY
  - Continue with 64 GB configuration
  - Number of VMs needs to be increased
    - January: 158
    - Near-term: no business forecast to use for sizing

How to improve capacity prediction and capacity efficiency

Capacity Risk highlighted

Risk: Usage needs to be **balanced** across VMs to make efficient use of configured capacity

APP C: RECOMMENDED VM CAPACITY

Capacity Prediction Risk highlighted

- CPU
  - **No business forecast**, so must use historical data
    - July Loc X: 82% utilization
      - 25 of 65 VMs had 4 instead of 8 vCPUs ➔ high utilization
    - Jan Loc X: 70% utilization
      - ½ the VMs over threshold
  - Sizing with transactions per VM
    - 64 VMs * 1.8 / .92 = 125 VMs
  - Sizing with users per VM
    - 4415 Users / 28 = 158 VMs
- MEMORY
  - **No business forecast**, so must use historical data
    - July Loc X: 40% utilization
    - Jan Loc X: 45% utilization
  - Loc Y config is the same as Loc X
    - Already meets threshold of 70%

Two sets of predictions shown; details in later slides

Analysis Risk: 2016 is running ahead of 2015, but there's **no business forecast** to explain this or to say what to expect in June/July

APP C: RECOMMENDED VM CAPACITY: USERS PER VM

- Users near 50 at peak workload, but server CPU above threshold of 70%
  - **Recommend using 28 users per VM for sizing**

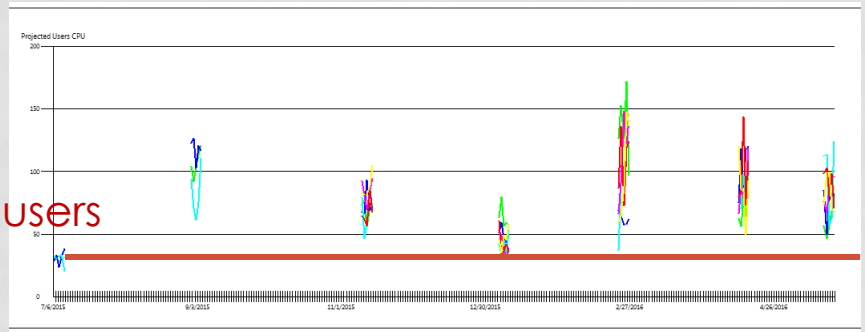

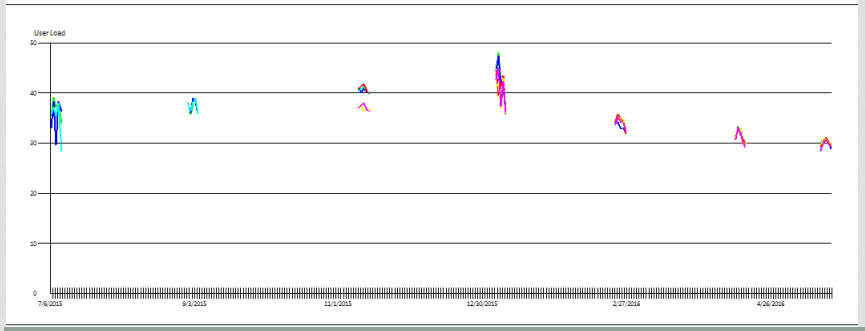

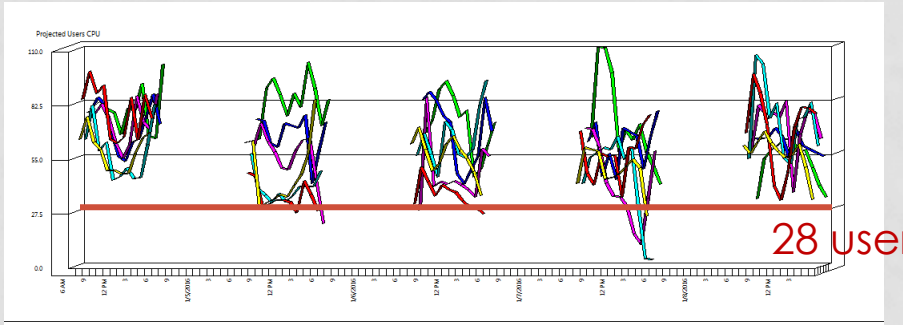

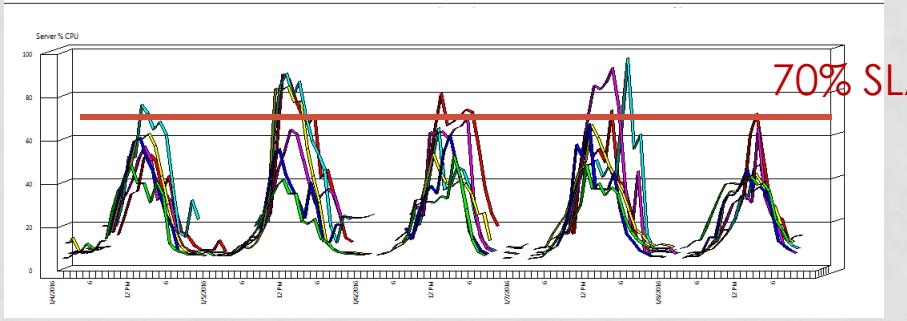

Aug 2015 – April 2016

Drill down for January 2016 peak

TECHNIQUE: PROJECTING CAPACITY USING UTILIZATION-BASED SLAS

- Two methods (pessimistic or optimistic)
  1. Assume total server/VM utilization represents the app
  2. Isolate app and calculate marginal cost of app load

Total CPU server utilization with SLA vs. App C CPU (dark blue)

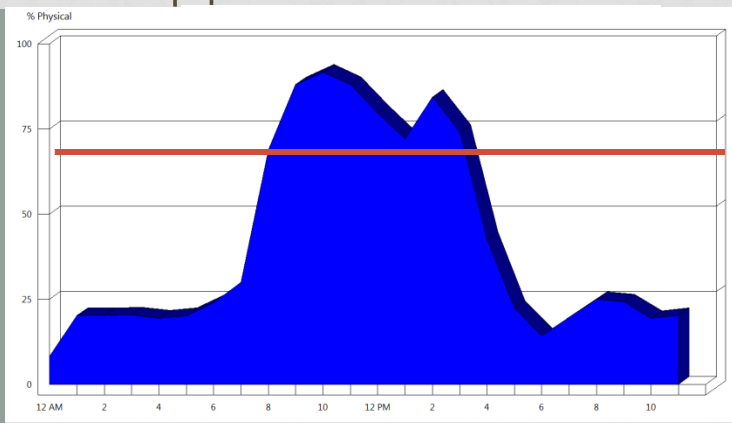

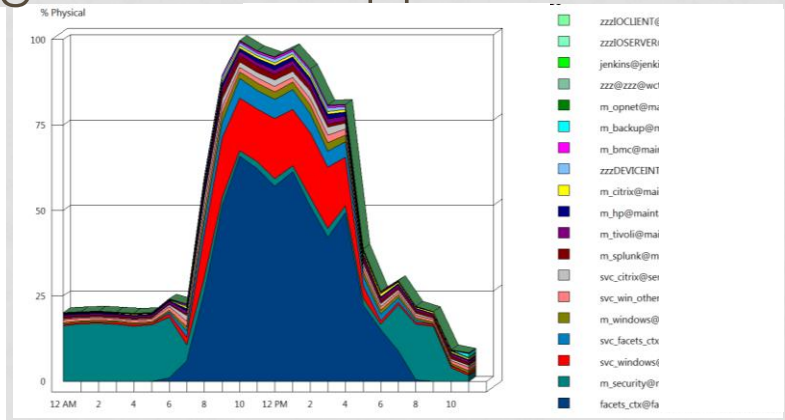

Observed number of users

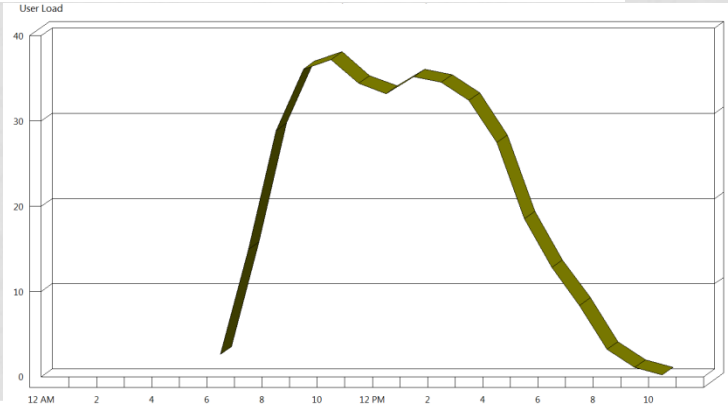

Cost per user =

App C CPU / users
or
Total CPU / users

TECHNIQUE: PROJECTING CAPACITY USING UTILIZATION-BASED SLAS

- Can be applied to any type of sizing where there's substantial "fixed costs", e.g. VMs per VMware host, application(s) on a server, etc.
- Use for any utilization-based SLA, e.g. CPU, Memory

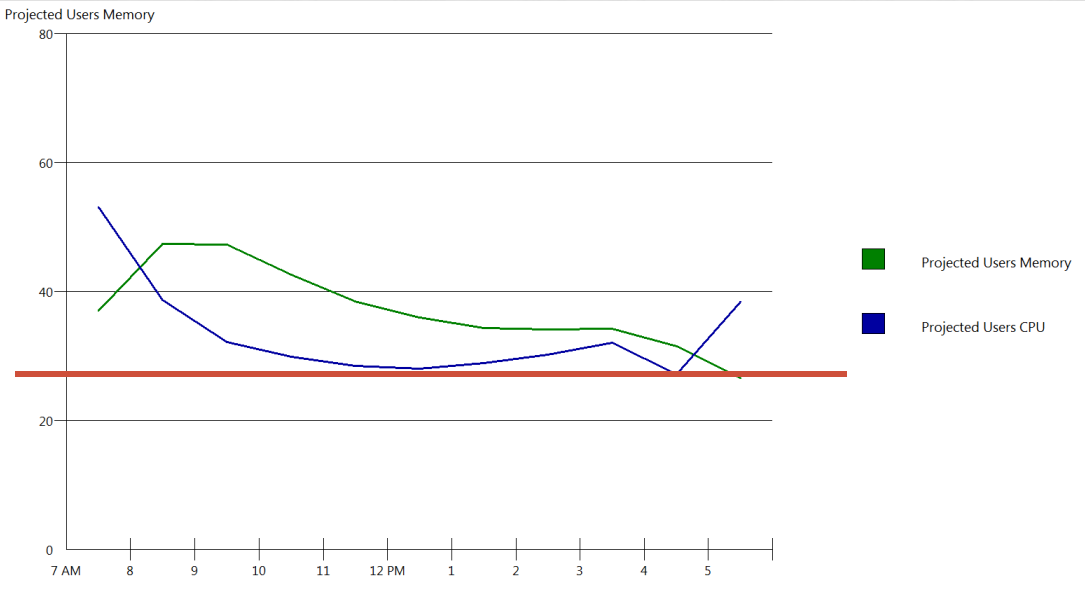

For capacity, take **worst case** during workload shift, e.g. 7 AM – 7 PM projects to 28 users
(or could use just peak workload period, 10 AM – 4 PM)

Projected user load = (CPU SLA – nonApp C CPU) / App C CPU per user

APP C: RECOMMENDED VM CAP
SPEC PER USER

- Historical CPU per user (selected VMs, selected v
  - Large variation by time of year, and across VM

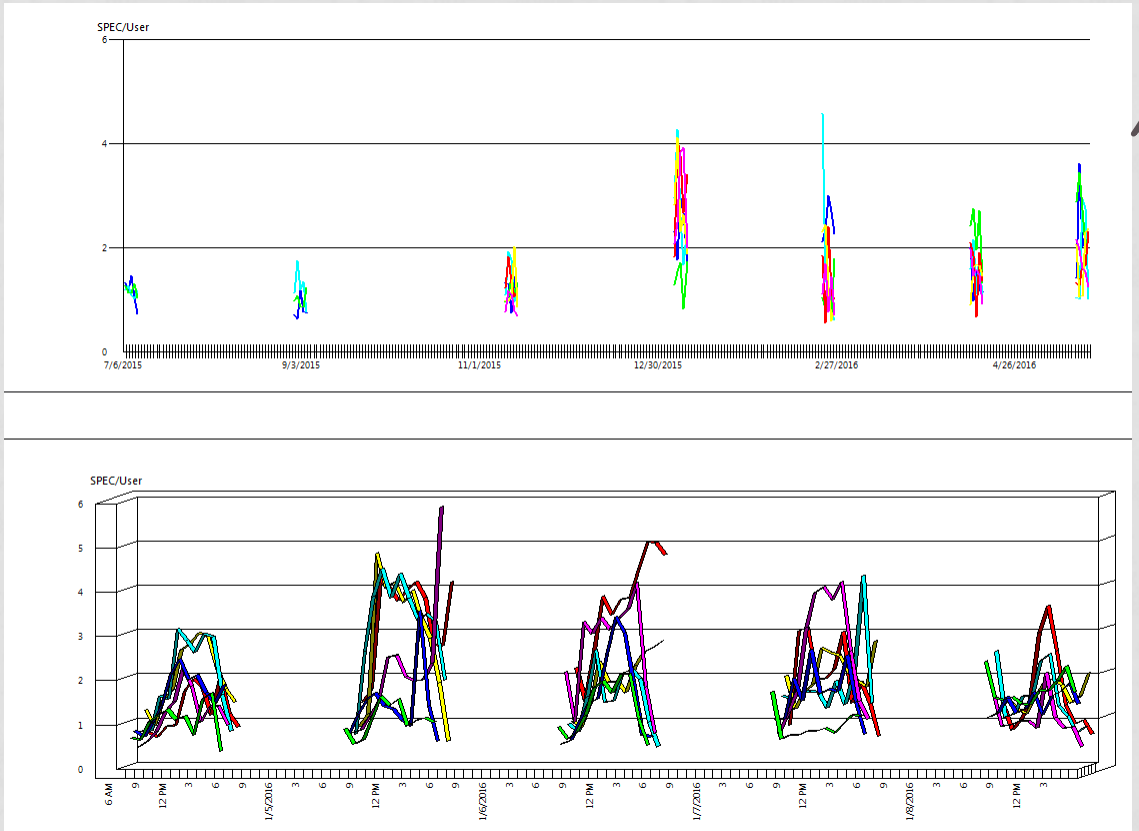

The marginal cost per user was isolated since server-level analysis would be too conservative

Aug 2015 – April 2016

Drill down for January 2016 business peak

Risk: Usage cannot be balanced because users don't do the same work

APP C: RECOMMENDED VM CAPACITY: TRANSACTIONS PER VM

- Compare observed daily transaction volume (total) with observed VM throughput to estimate number of VMs needed on that day to support the volume
  - Shown with threshold of currently configured 96 VMs

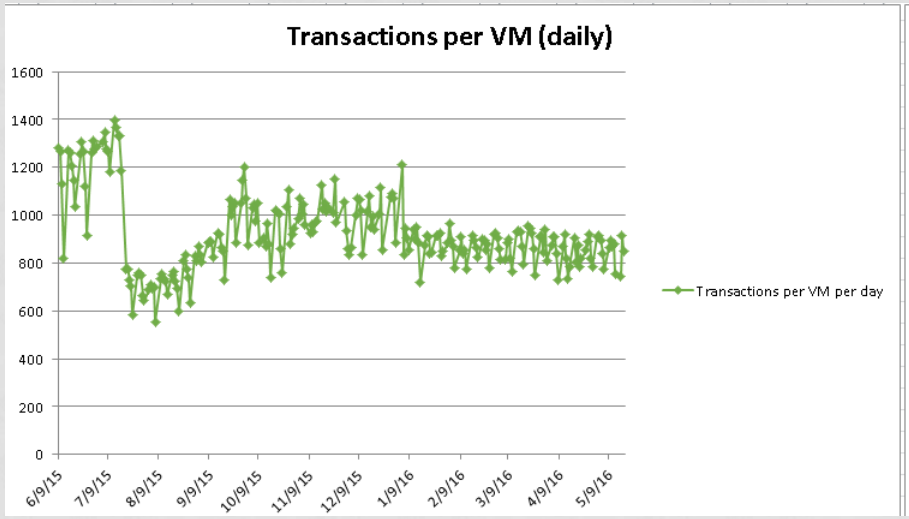

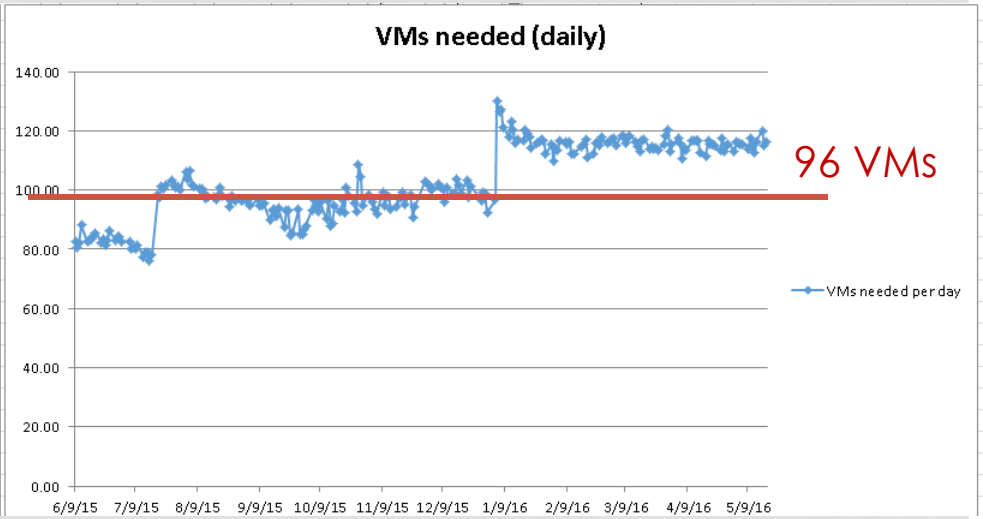

June 2015 – May 2016

Note: Not adjusted for 70% CPU threshold: January requirement would be higher, April-May would be lower

TECHNIQUE: IGNORE BAD CAPACITY RESULTS DUE TO OPERATIONAL PROBLEMS

- Always review "most important" data points for problems/typicality
  - Identify specific low (or high) results to verify that most important elements are intact
    - For this study, that would be server CPU utilization, APP C CPU utilization, and non-App C CPU

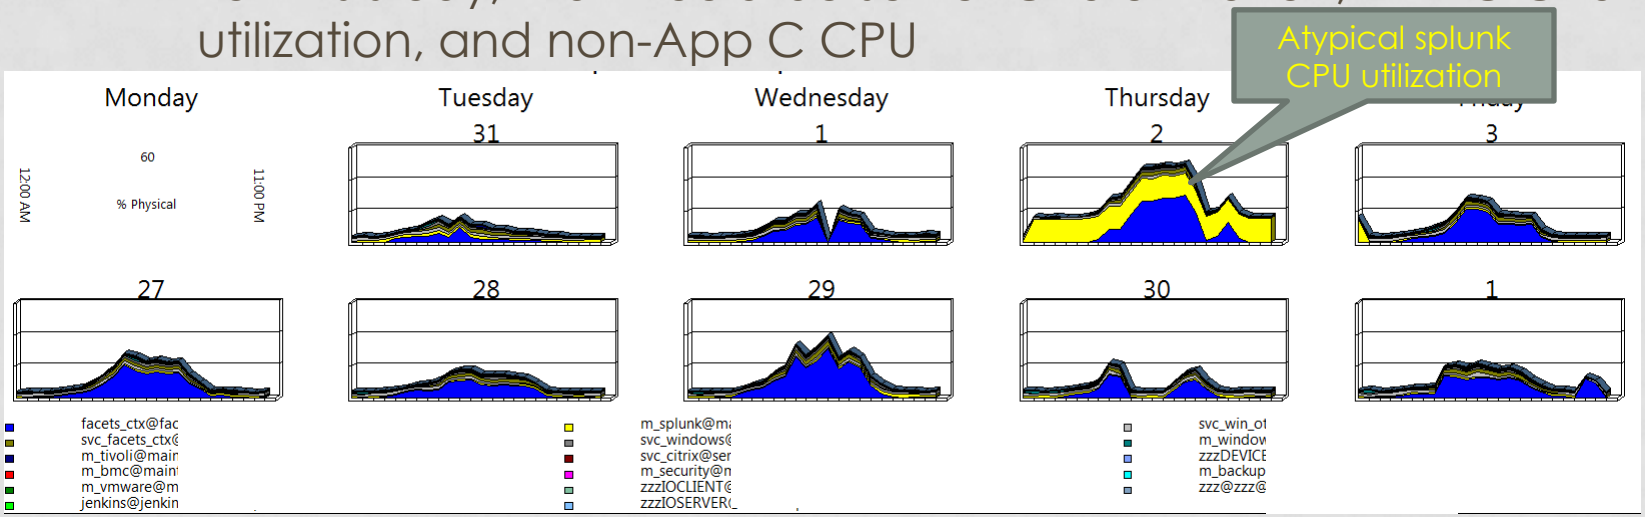

Don't use any results from this server from Thursday the 2nd since splunk needed to be recycled!