Virtualization Management Index

Issues 1 and 2

Virtualization Management Index

Issue 1

December 2010

REPORT BY BRYAN SEMPLE

Key Findings

The Virtualization Management Index (VMI) is the virtualization industry's first empirical study of virtualized data centers. Drawn from a data set size of over 2,500 unique environments and 550,000 virtual machines, the VMI provides a glimpse into the status of server consolidation measured by by number of virtual machines per host organizations are able to attain. Today, the VMI stands at 15.6 VMs/host with memory per virtual machine the primary limiting factor for additional consolidation. Enterprises should perform exhaustive TCO analysis to determine if the cost of extreme memory density configurations offsets the cost of added servers and respective software licensing.

Environment Statistics

VMI data is collected from some of VKernel's free tools providing an interesting snapshot into the resource efficiency of today's virtualized environments. Inclusion in the VMI was limited to environments of more than 50 virtual machines. Data collected includes:

- Hosts
- VMs
- Clusters
- Resource Pools
- Storage (allocated, attached)
- Memory (allocated, available)
- CPU (allocated, available)
- Power On, Power Off VMs
- Cores, Sockets, vCPU
- VMs with performance problems – CPU, memory, storage, latency
- VMs underutilized – CPU, memory, storage

The average environment size was 225 virtual machines and 18 hosts. An average host consisted of:

- dual socket (2.4 sockets)
- quad core (3.6 cores) servers
- 50GB of memory
- 2.6GHZ on CPU

Virtual resource usage was:

- 0.9 VMs/vCPU
- 2.2 vCPU/core
- 2 VMs/core
- 2.25 GB memory allocated/VM

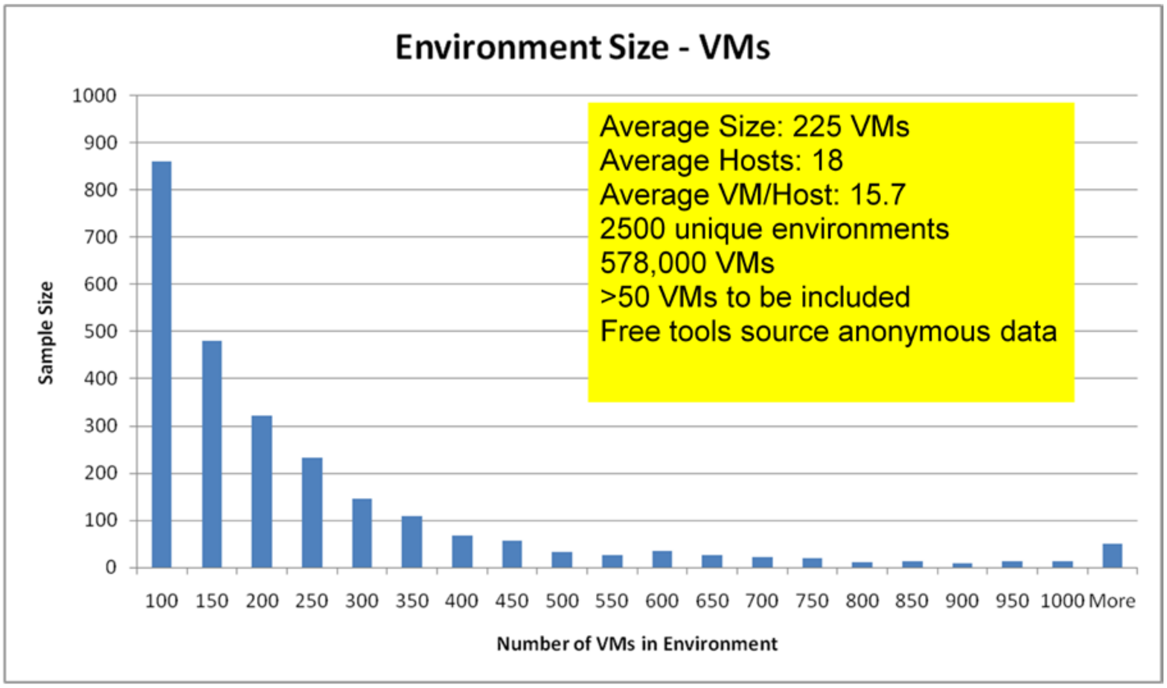

Consolidation Efficiency: 15.6 VMs/host

The VMI shows a consolidation efficiency of 15.6 VMs/host. The consolidation efficiency appears to vary based on environment size with larger environments being less efficient at VM deployment.

Size of Environment (Hosts)	VM/Host Ratio
<10	20
10 – 24	13
>24	11
Average by environment	15.6
Average by sample	12.5

From examining both CPU and memory allocation information, memory is the limiting factor to more consolidation. The ratio of allocated to physical memory is 0.7. System administrators are only comfortable allocating 70% of the memory and keeping the remaining 30% in reserve.

CPU allocation is another story. The ratio of allocated to actual CPU cycles is 2.2. System administrators are very comfortable allocating more CPUs cycles than are physically available due to the very low CPU utilization of most applications.

Hence memory is the constraining resource. Increasing memory density involves a tradeoff between higher cost, higher density memory sticks and lower density memory that yields a lower total memory capacity for a server. CPU density is relatively cheap compared to memory. So the servers in the survey represented the typical case where environments have plentiful cores, but not sufficient memory. This also results in a relatively low VM per core ratio of 2.2.

Further investigation should be conducted to determine where the breakeven point exists for higher density memory. Doubling memory density from 16 to 32 GB, for example, usually involves more than a 2x cost increase. However, at some point, the ability to increase CPU utilization higher will make this increased investment worthwhile by deferring incremental server purchases and the resulting operating and licensing expenses.

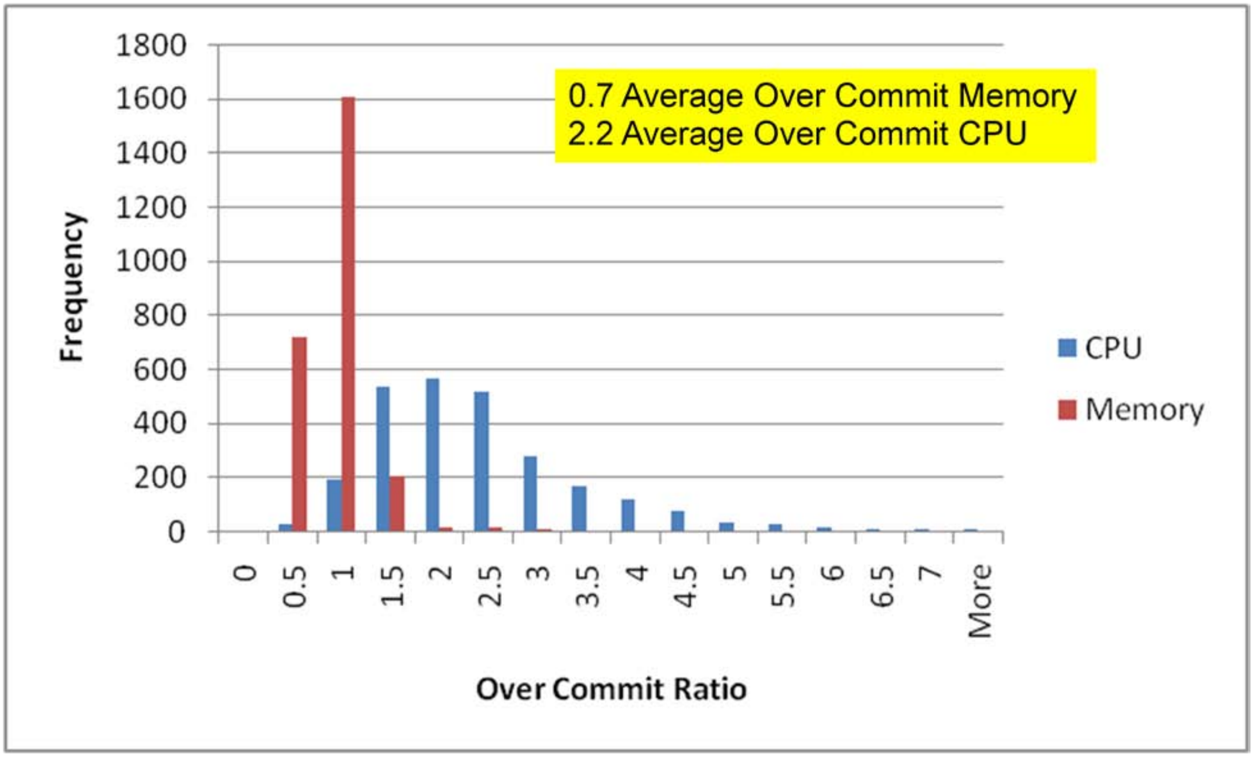

Density Outliers

While 0.9 VMs/vCPU and 15.6 VMs/host was the average for the index, some environments safely increased VM density. By finding environments with allocated to actual CPU and memory ratios greater than 1, the original 2,500 data centers were reduced by 90%. Data centers that had no performance issues were then mined from the remaining sample size yielding 92 data centers or 3.7% of the total. These data centers had extremely noteworthy characteristics. They achieved an average density of 25.5 VMs/host and an average cost that was 50% lower than the average while not triggering any performance alerts.

Metric	High Density	Average
VMs/Host	25.5	15.6
VMs/vCPU	1.1	.93
vCPU/Core	3.3	2.17
CPU Allocation Ratio	3.3	2.18
Memory Allocation Ratio	1.2	.66
CPU and Memory Bottlenecks	0%	.3%

VMI Statistics Round Up

Environment Statistics	Average	Median	Max
Hosts	18.1	11.4	177.2
Clusters	3.7	2.0	50.8
Data Stores	53.5	33.2	1308
Total VMs	225.6	142.0	2,955
Sockets	43.1	26.0	510
Cores	153.3	92.5	2,058
Memory (GB)	897	496	14.9 PB
Storage(TB)	28	16.8	525

Per VM, Per Host	Average	Median	Max
VMs/ Host	15.7	13.0	200
VMs/Core	1.9	1.6	17.5
Sockets/Host	2.4	2.0	8.0
Cores/Socket	3.6	4.0	7.0
Memory/Host	50.0	41.4	
Storage/VM	85.5	65.1	
Storage/Host	1.8 TB	1.4 TB	25.6 TB
Memory/VM (GB)	4.0	3.4	28.5
Hosts/Cluster	5.7	4.0	

Conclusion

The VMI yields a glimpse into virtualization. Memory density continues to hamper increases in VM density. Enterprises should perform extensive TCO analysis to determine if extreme memory densities are a worthwhile investment to increase VM density and lower server count.

Virtualization Management Index

Issue 2

January 2011

REPORT BY BRYAN SEMPLE

Key Findings

The Virtualization Management Index (VMI) is the virtualization industry's first empirical study of virtualized data centers. This issue is an update to the first index and confirms that larger environments have lower VM densities than smaller environments resulting in approximately 50% greater cost per VM for these larger enterprises.

Trimming the Data Set

Our first VMI issued in December 2010 contained information that seemed to point to the fact that larger environments had worse VM densities than smaller environments.

Size of Environment (Hosts)	VM/Host Ratio
<10	20
10 – 24	13
>24	11
Average by environment	15.6
Average by sample	12.5

This led to a flurry of speculation from bloggers and analysts:

- VDI environments are skewing the data
- Powered off VMs are skewing the data
- Larger environments have older equipment which supports fewer VMs

To provide further clarity, the original data set was trimmed down further with the following requirements:

- Only environments with more than 50 VMs and more than 2 hosts were considered to eliminate even more potential test and dev systems.
- No more than 15% of the environment could be powered off to remove VDI and more test and dev
- Only environments with at least one cluster were considered to remove smaller environments from the sample that were potentially not production

The result was a reduction in the sample size from 2,500 unique environments to 1,528 environments.

Large Environments Have Poor VM Densities

Despite removing environments with high numbers of powered off VMs and smaller environments without sufficient hosts, the results for VMs/host did not vary dramatically from the original study.

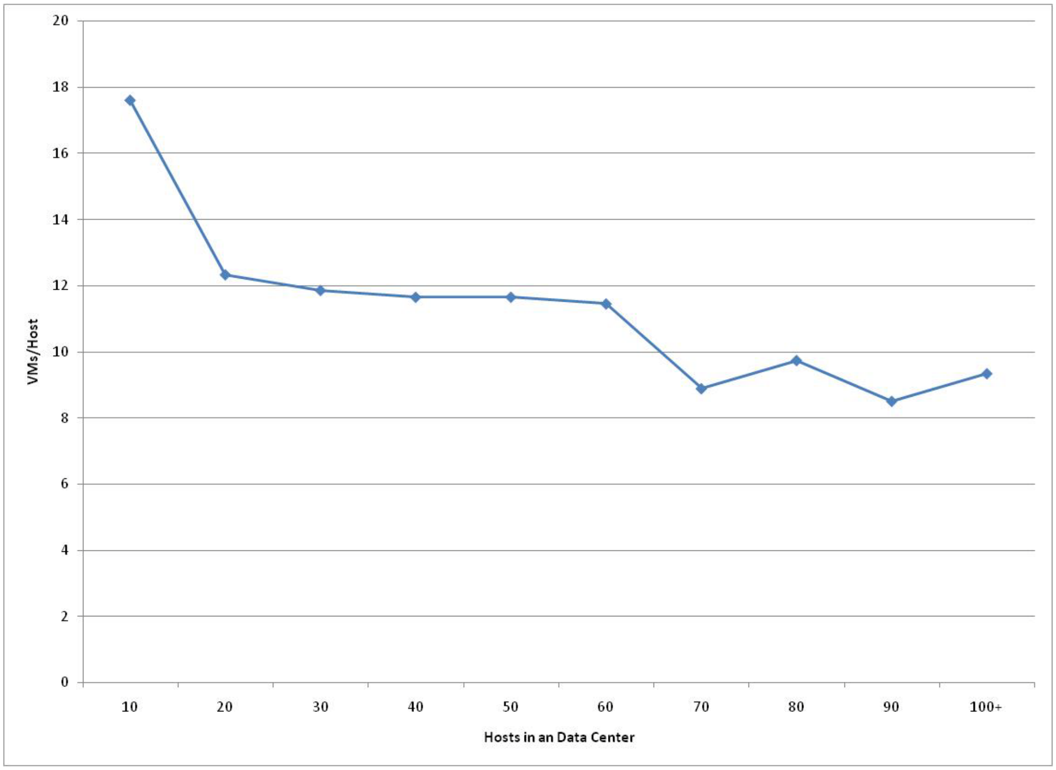

Smaller environments with 10 or less hosts had the best VM/host density approaching nearly 18. As environments grew in size, there was a dramatic reduction in VM density. The biggest reduction occurred as the host count increased to 20 hosts with a gradual decrease occurring as host count increased above 100.

There was no statistical difference in cluster size, powered on VMs, or hardware capabilities between the host size groups. There are two apparent conclusions:

- As environments grow larger, system administrators are unable to properly manage consolidation ratio's to keep virtualized environment efficiency high
- Performance uncertainty as environments scales forces significant over allocation of resources to compensate for performance uncertainty.

Implications for Directors of Infrastructure

Essentially, larger environments are paying almost 50% more per VM than smaller environments. The cost data is drawn from the original VMI study. Why do these organizations tolerate the high cost? Possible causes could be:

- There are still significant savings to be gained from virtualizing physical applications hence there is little pressure to be more efficient with existing virtualized servers.
- Performing capacity management on large scale environments is nearly impossible without analytics. The result is tremendous uncertainty on actual capacity needs. To compensate for this

uncertainty, over allocation of hardware resources is used. While theoretically effective at reducing performance problems, the data reveals the rate of memory and CPU induced performance problems is independent of VM density ratios.

- Uncertainty over capacity needs will hamper deploying mission critical applications. Mission critical applications require known resource requirements to prevent performance bottlenecks.

Conclusion

The first VMI provided a glimpse into the status of virtualization including the observation that larger environments appeared to be less efficient with VM densities. This update reinforced the observation that larger virtualized environments suffer from poorer VM density than smaller environments and hence have up to 50% higher cost per VM. Directors of infrastructure will need to confront this issue as the low hanging fruit of cost reduction by virtualizing physical applications is removed and CIOs look for the next wave of savings. The savings can be achieved, but only through analytical capacity management.