

Testing RESOURCE POOL Settings

The purpose of this test case is to explore the performance impact of various RP (Resource Pool) Settings and to understand what each RP setting means in regards to setting reservation and limit amounts for CPU and Memory.

Test Environment: In the Test VMware environment in the IOlab we are working off vCenter 4.0 server IODWMGT002, (2) ESX 4.0 hosts in a DRS Cluster, and with (6) Windows 2003 x86 VMs running various CPU and Memory burner applications to generate loads under various RP configurations. Specifically we are exploring the impact of setting the RP limit value to static setting vs. unlimited in a fixed resource pool (Expandable NOT checked).

General	
Manufacturer:	HP
Model:	ProLiant BL460c G1
CPU Cores:	8 CPUs x 2.333 GHz
Processor Type:	Intel(R) Xeon(R) CPU E5345 @ 2.33GHz

Environment Constants:

No VM level reservations/limits configured.

No Share values configured anywhere.

Fixed Resource Pools ("Expandable Reservation" NOT checked)

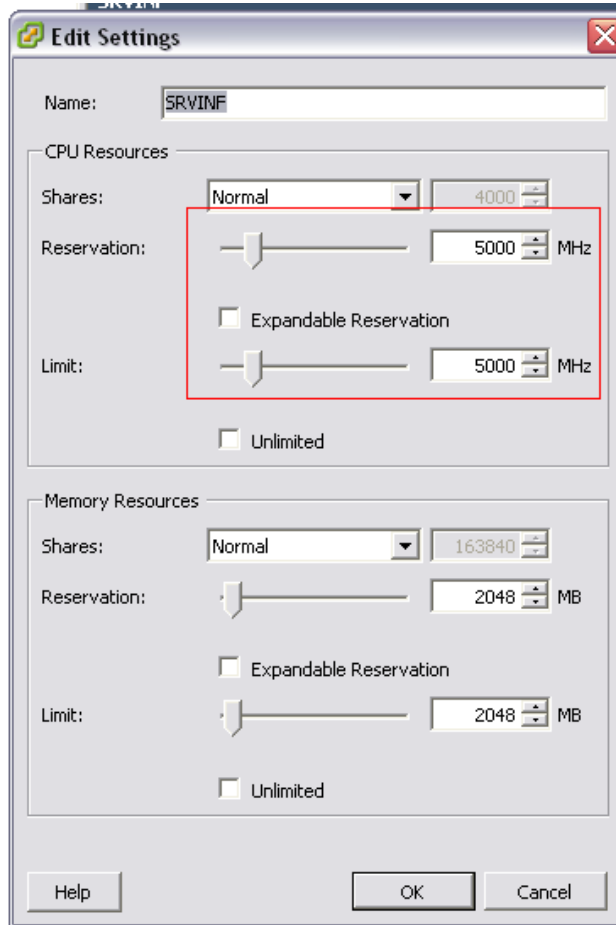

Environment Variables:

CPU and Memory Burner tools run in each of (4) test VMs. (ENG_TEST_2003_XX)

Limit is changed from matching the reservation amount to unlimited.

CPU TESTING-----

With Limits and No load:



CPU

Host CPU

0 MHz 5000 MHz

Consumed 0 MHz
Active 0 MHz

Resource Settings

- ▾ Reservation 5.00 GHz
- ▾ Limit 5.00 GHz

Shares Normal (4000)
Worst Case Allocation 5.00 GHz

[Help](#) Edit

SRVINF

Summary Virtual Machines **Resource Allocation** Performance Tasks & Events Alarms Permissions Maps Storage Views

CPU

Configured Reservation: **5000 MHz**

Reservation Type: **Fixed**

Used Reservation: **0 MHz**

Available Reservation: **5000 MHz**

Memory

Configured Reservation: **2048 MB**

Reservation Type: **Fixed**

Used Reservation: **804 MB**

Available Reservation: **1244 MB**

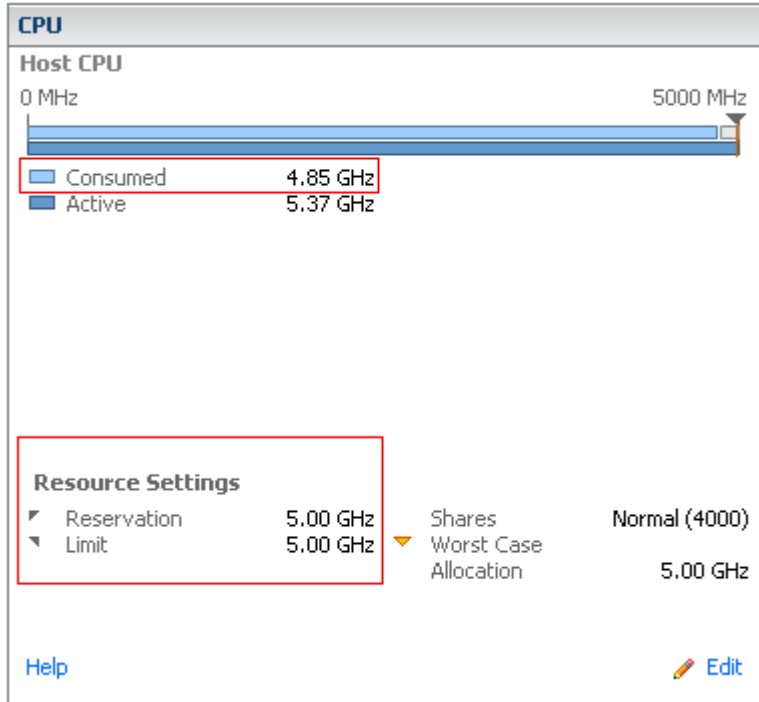
View: [CPU](#) [Memory](#)

Name	Reservation - MHz	Limit - MHz	Shares	Shares Value	% Shares	Worst Case Allocation - MHz
ENG_Test_2003_10	0	Unlimited	Normal	1000	16	833
Test001	0	Unlimited	Normal	1000	16	835
Test002	0	Unlimited	Normal	1000	16	833
ENG_Test_2003_11	0	Unlimited	Normal	1000	16	833
ENG_Test_2003_12	0	Unlimited	Normal	1000	16	833
ENG_Test_2003_13	0	Unlimited	Normal	1000	16	833

[Edit SRVINF resou](#)

*From the screenshots we can see at idle these VMs will share the available CPU reservation equally.

With Limits and CPU load:



SRVINF

Summary Virtual Machines Resource Allocation Performance Tasks & Events Alarms Permissions Maps Storage Views

Name, State, Host or Guest OS contains: Clear

Name	State	Status	Host	Provisioned Space	Used Space	Host CPU - MHz
Test001	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	0
Test002	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	0
ENG_Test_2003_10	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	1236
ENG_Test_2003_11	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	1213
ENG_Test_2003_12	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	1189
ENG_Test_2003_13	Powered On	Normal	iodvhst006.dev.sutterhealth.org	21.00 GB	21.00 GB	1143

*From the images below we can see that when the test machines (ENG_TEST_2003_XX) run a CPU burner they will choke out the other VMs in the RP and will evenly take the available CPU reservation.

Unlimited and CPU load:

Edit Settings [Close]

Name:

CPU Resources

Shares:

Reservation: MHz

Expandable Reservation

Limit: MHz

Unlimited

Memory Resources

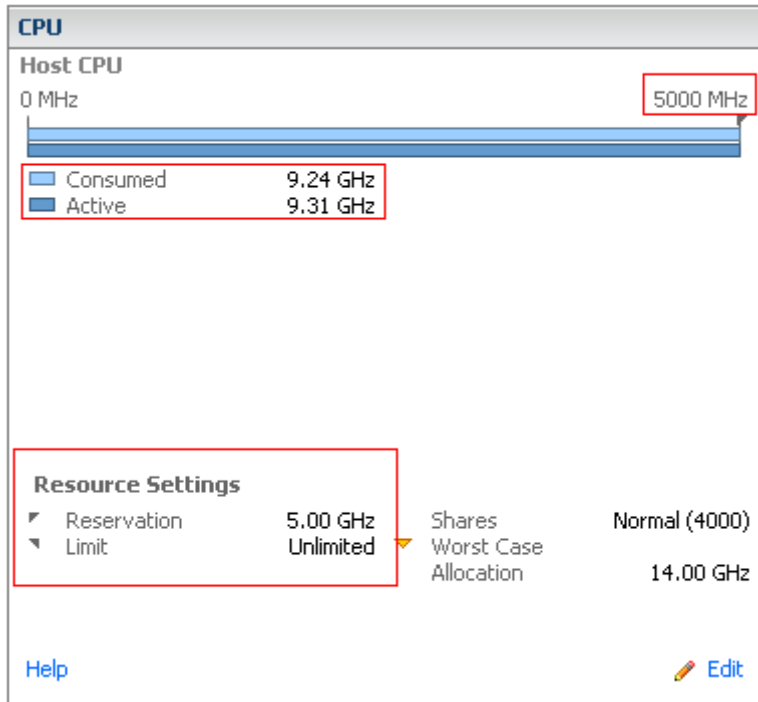
Shares:

Reservation: MB

Expandable Reservation

Limit: MB

Unlimited



VMs are now taking more and more CPU at an unlimited rate beyond the reservation.

Name	State	Status	Host	Provisioned Space	Used Space	Host CPU - MHz
Test001	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	0
Test002	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	0
ENG_Test_2003_10	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	2309
ENG_Test_2003_11	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	2309
ENG_Test_2003_12	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	2309
ENG_Test_2003_13	Powered On	Normal	iodvhst006.dev.sutterhealth.org	21.00 GB	21.00 GB	2333

*From the images above we see that will limit set to “unlimited” the VMs will take up as many resources from the parent as needed.

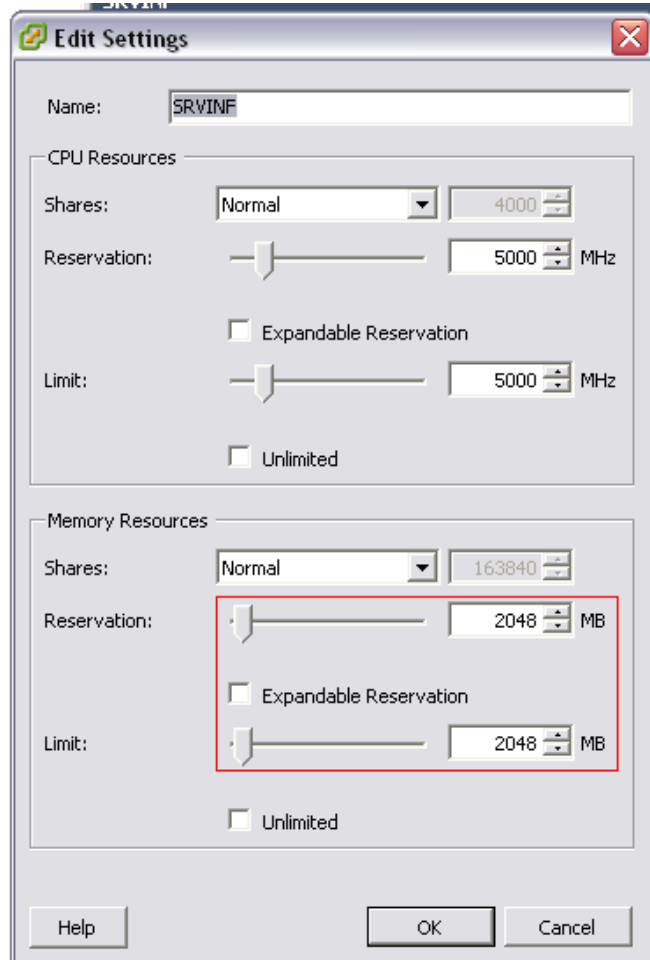

Conclusion: Each VM with a single CPU can potentially use up to 2300 Mhz of the host CPU. This is equal to speed of a single core. This means each single vCPU VM can use up to a single core.

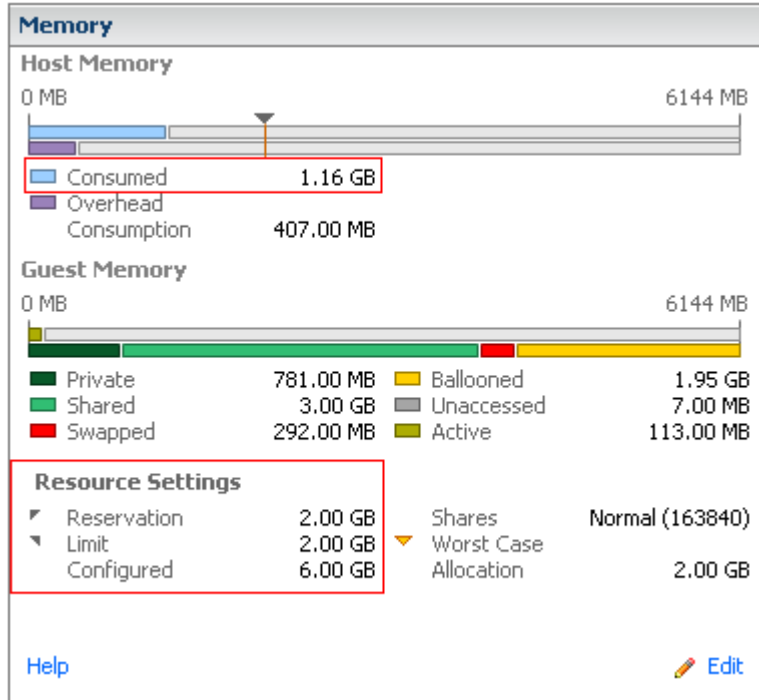
When a limit value is set this is an aggregate cap on the max amount of CPU all the VMs in a RP can use. If a reservation limit is set and the limit value is higher or unlimited, then the delta represents the amount of CPU resources that the VMs can attempt to take from the parent RP (This is the cluster itself in our case).

When VMs in the same RP compete for resources in their pool or the parent, they each have equal share values so they divide up resources equally.

Memory TESTING-----

With Limits and No load:





SRVINF

Summary Virtual Machines Resource Allocation Performance Tasks & Events Alarms Permissions Maps Storage Views

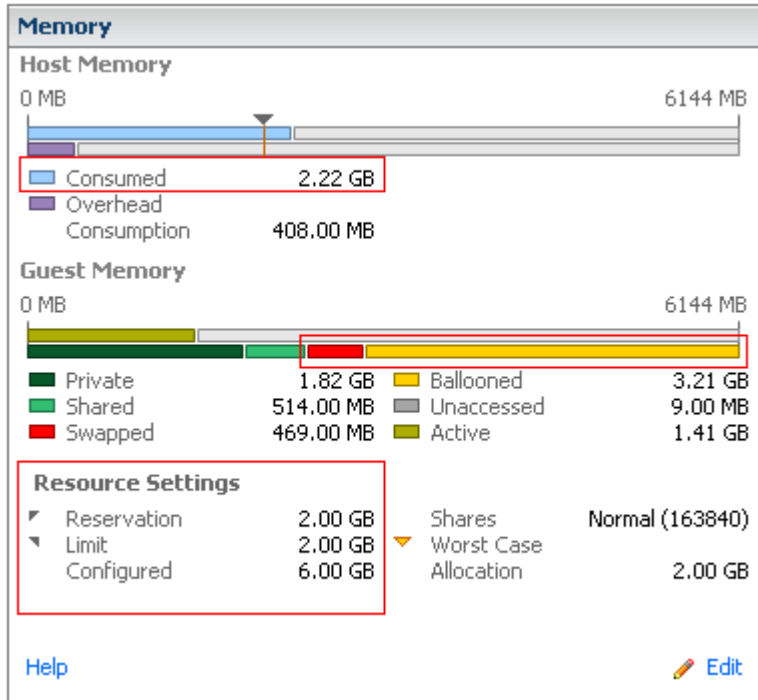
CPU		Memory	
Configured Reservation:	5000 MHz	Configured Reservation:	2048 MB
Reservation Type:	Fixed	Reservation Type:	Fixed
Used Reservation:	0 MHz	Used Reservation:	804 MB
Available Reservation:	5000 MHz	Available Reservation:	1244 MB

View: CPU Memory Edit SRV

Name	Reservation - MB	Limit - MB	Shares	Shares Value	% Shares	Worst Case Allocati...
ENG_Test_2003_10	0	Unlimited	Normal	10240	16	341
Test001	0	Unlimited	Normal	10240	16	341
Test002	0	Unlimited	Normal	10240	16	341
ENG_Test_2003_11	0	Unlimited	Normal	10240	16	343
ENG_Test_2003_12	0	Unlimited	Normal	10240	16	341
ENG_Test_2003_13	0	Unlimited	Normal	10240	16	341

*The above images show that at idle VMs will take an even amount of the available reserved memory.

With Limits and Memory load:

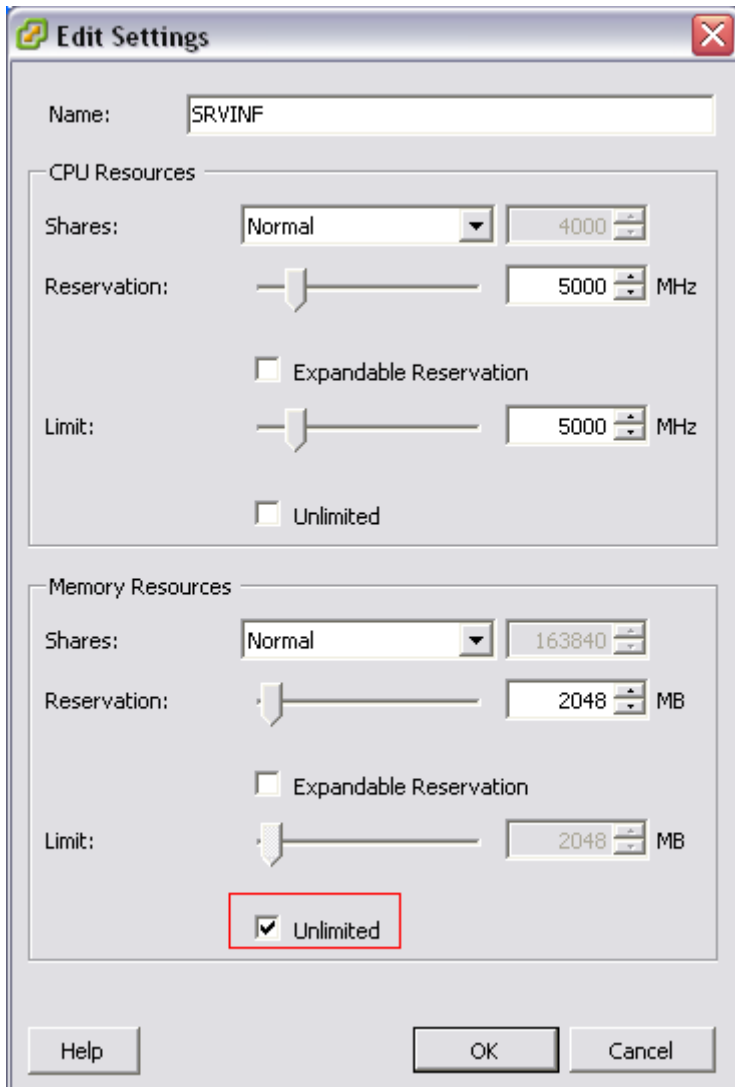


*Because the reservation is below the configured amount and the VMs are being benchmarked we are seeing lots of Ballooning and Swapping. –The normal signs that VM memory is constrained.

Host	Provisioned Space	Used Space	Host CPU - MHz	Host Mem - MB	Guest Mem - %
iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	0	123	2
iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	0	165	0
iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	1283	713	43
iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	1283	519	43
iodvhst006.dev.sutterhealth.org	21.00 GB	21.00 GB	1143	424	23
iodvhst006.dev.sutterhealth.org	21.00 GB	21.00 GB	1143	335	26

*The above images show that under a memory load that stressed VMs will take up the available memory evenly up to the reserved amount.

Unlimited and Memory Load:



SRVINP									
Summary Virtual Machines Resource Allocation Performance Tasks & Events Alarms Permissions Maps Storage Views									
Name, State, Host or Guest OS contains: []									
Name	State	Status	Host	Provisioned	Used Spa...	Host CPU - MHz	Host Mem - MB	Guest Mem - %	
Test001	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	0	127	2	
Test002	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	0	167	2	
ENG_Test_2003_10	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	2333	985	73	
ENG_Test_2003_11	Powered On	Normal	iodvhst005.dev.sutterhealth.org	21.00 GB	21.00 GB	2309	989	89	
ENG_Test_2003_12	Powered On	Normal	iodvhst006.dev.sutterhealth.org	21.00 GB	21.00 GB	2309	1000	81	
ENG_Test_2003_13	Powered On	Normal	iodvhst006.dev.sutterhealth.org	21.00 GB	21.00 GB	2309	986	84	

*The above images show that memory stressed VMs in an unlimited RP will use the reservation and then take memory from the parent until each VM reaches its configured amount.

Conclusion: Each VM has the ability to use up to the configured amount per VM. How this configured amount translates to actual resource usage depends on the RP setting. In a fixed RP with a defined limit

below the aggregate of the configured amounts, there will always be contention. If the pool is set to unlimited, VMs will take as much memory they need up to their configured amount. This could mean using all of the reservation in the pool and then going to the parent and taking what is needed up to the configured amount per VM.

FINAL CONCLUSION-----

The limit value is a very important resource pool setting that dictates the hard limit cap of how much CPU or Memory will be available to all the VMs in a RP. This hard limit can be reached in a number of ways depending on whether we are using Expandable Reservation (fixed/not fixed) or if we are using a defined limit or unlimited.

It all boils down to a math game of comparing the RP's defined reservation amount VS fixed/not fixed VS limit/no limit VS the configured amount of resources per VM. This is a game where we must decide if we want to over allocate or not over allocate.

CPU: When each VM gets a single vCPU, if the RP is unlimited each VM could use up to an entire core worth of CPU. On a 2-way Quad Core ESX host this means 8 runaway VMs could use all the CPU. With that said it is important to NOT use "Unlimited" for CPU on Resource Pools.

Memory: Because each VM's max possible memory usage is defined by the per VM configured amount, it's important that the RP limit amount is the same as the aggregate of all child VMs configured memory amount. If the reservation meets this amount (limit) or less then that is a decision on how much to oversubscribe memory.