





CARD 1

Card 2

SSD RAID 1
M2 S21

RAID 1 across Fusion IO cards




300GB 15K

7 Disk RAID5 7 Disk RAID5


1.8TB

1.8TB


52" SPANNED



Windows Spanned volume using Dynamic Disks. Writes only happen to Disk 4 until it is completely full and then the writes span over to Disk 5, meaning 50% of IO potential is wasted in this setup. Rebuild these to RAID 10 or other configuration.



AMD Magny-Cours, Interlagos, Bulldozer, and Piledriver processor configuration for NUMA. Each of these represents 2 NUMA nodes and has 2 memory controllers per socket.



NUMA connection assignment – round robin across NUMA nodes for the default TCP Listener port, with task assignment being made to the least load_factor scheduler in the node as the initial ideal scheduler. If the ideal scheduler is 120% higher load than another scheduler in the node upon new task assignment, it will change the ideal scheduler to the scheduler with the least load_factor to balance the load across schedulers in a single NUMA node.

Max Server Memory

Slide 17

1GB OS

3GB 4-16GB

6GB 16-64GB

64GB

1

12/4

48/8

64

- 12

54GB

1GB

Example scenario using the 'max server memory' calculation: 1GB for OS + 1GB for each 4GB between 4-16GB + 1GB for each 8GB > 16GB using a server with 64GB memory. The resulting configuration would be 'max server memory' set at 54GB initially then monitor Memory\Available Mbytes and increase based on the lowest memory point in 24 hr periods by no more than 512MB – 1GB at a time.


Max Transfer Size

4MB

Backup Memory

Buffer Count

1000




4GB Backup Memory

4MB

x64 Buffers

256MB



Mem To Leave

Backup Buffer memory usage is $\text{MAXTRANSFERSIZE} * \text{BUFFERCOUNT}$ and was problematic on 32-bit instances due to the small (~384MB) MEMTOLIEAVE