Re: VFS sysctl tuning

Source: http://unix.derkeiler.com/Mailing−Lists/FreeBSD/performance/2006−08/msg00010.html

- *From*: "Nikolas Britton" <nikolas.britton@xxxxxxxxx>
- *Date*: Sat, 12 Aug 2006 22:37:33 −0500

On 8/12/06, Nikolas Britton <nikolas.britton@xxxxxxxxx> wrote:

> On 8/12/06, Robert Watson <rwatson@xxxxxxxxxxx> wrote:
> > On Sat, 12 Aug 2006, Nikolas Britton wrote:
> >
> > > After reading man tuning it said to change vfs.hirunningspace to a larger
> > > value. What should I set it to?
> > >
> > > My RAID controller, an Areca ARC−1220, has a 256MB cache and each disk has
> > > 16MB of cache.
> > >
> > > What other vfs.* sysctls can I change?
> >
> > The default settings are "reasonable" for many common workloads. Further
> > tuning generally should be done in the context of a specific workload, since
> > tuning generally means trading off one type of performance or functionality
> > for another. Only by knowing the workload can we talk about how to trade off
> > what you need, and what you don't.
> >
>
> Thanks Robert. It's a general purpose server but It's primary duty is
> serving files via Samba / NFS over GigE to workstations that are
> faster then it. We are typically working with hundreds of (new) ~16MB
> files on any given day. It currently has a single dual−core Xeon 5030
> and 1GB of system ram. I plan on upgrading it to dual dual−core Xeon
> 51xx and maxing out the ram if FreeBSD ever gets Xen dom0 support so
> we can consolidate server equipment... Anyhow... I decided to script
> out the task of finding the right setting for the hirunningspace
> sysctl:
>
> #!/bin/csh
> sysctl vfs.hirunningspace=1048576
> iozone −a −g 1g −b ~/arc1220_18.xls
> sysctl vfs.hirunningspace=2097152
> iozone −a −g 1g −b ~/arc1220_19.xls
> sysctl vfs.hirunningspace=4194304
> iozone −a −g 1g −b ~/arc1220_20.xls
> sysctl vfs.hirunningspace=8388608
> iozone −a −g 1g −b ~/arc1220_21.xls

sysctl vfs.hirunningspace=16777216
iozone −a −g 1g −b ~/arc1220_22.xls
sysctl vfs.hirunningspace=33554432
iozone −a −g 1g −b ~/arc1220_23.xls
sysctl vfs.hirunningspace=67108864
iozone −a −g 1g −b ~/arc1220_24.xls
sysctl vfs.hirunningspace=134217728
iozone −a −g 1g −b ~/arc1220_25.xls
sysctl vfs.hirunningspace=268435456
iozone −a −g 1g −b ~/arc1220_26.xls

Ok first I should say that the default 1MB setting is pretty darn
good. It had the highest composite index score out of all the tests
BUT the 32MB setting had the second best composite index score and had
substantially better write scores in the larger test file sizes.

Here are the iozone write test averages:
column 1 = hirunningspace=1MB
column 2 = hirunningspace=32MB
column 3 = Size of test file.
all numbers are in KB/s

310683 vs. 267439: 64
284124 vs. 251948: 128
235227 vs. 214273: 256
260533 vs. 235084: 512
257604 vs. 245767: 1024
288502 vs. 223431: 2048
297675 vs. 269663: 4096
299051 vs. 302620: 8192
304752 vs. 276025: 16384
322887 vs. 323967: 32768
287164 vs. 272954: 65536
222640 vs. 255020: 131072
201985 vs. 226721: 262144
188920 vs. 214589: 524288
190022 vs. 216239: 1048576

Now I'm running these tests (24MB, 28MB, 40MB, 48MB):
sysctl vfs.hirunningspace=25165824
iozone −a −g 1g −b ~/arc1220_27.xls
sysctl vfs.hirunningspace=29360128
iozone −a −g 1g −b ~/arc1220_28.xls
sysctl vfs.hirunningspace=41943040
iozone −a −g 1g −b ~/arc1220_29.xls
sysctl vfs.hirunningspace=50331648
iozone −a −g 1g −b ~/arc1220_30.xls

I have a hunch that the best results are going to be 1/4 of

vfs.maxbufspace, which would be ~28MB. We'll see.

−−
BSD Podcasts @:
http://bsdtalk.blogspot.com/
http://freebsdforall.blogspot.com/