
Subject: Re: Re: TCP packets lost?

Posted by [gblond](#) on Mon, 07 Jan 2008 14:43:48 GMT

[View Forum Message](#) <> [Reply to Message](#)

On 7 January 2008 15:02:14 Max Deineko wrote:

> Vitaliy Gusev <vgusev@sw.ru> wrote:

> > On 7 January 2008 13:01:24 Max Deineko wrote:

> >> > On Mon, Jan 07, 2008 at 07:55:43AM +0000, Max Deineko wrote:

> >> >> My first guess would have been that when the server is busy

> >> >> serving an expensive request, the tcp rcv buffer would get full

> >> >> and packets would get dropped. But wouldn't tcpdump not be able

> >> >> to see them then?

>

> >> I already [...] tried increasing tcp_rmem size, which hung the

> >> webserver completely (maybe it was a litte too big ;), but I'm still

> >> experimenting with the settings.

> >>

> >> Of course there might be a problem with other server components, and

> >> I'm not really a networking expert, it's just that the tcp flow

> >> doesn't look like what I'd expect it to.

>

> > Please show files: /proc/net/netstat, /proc/net/snmp (inside VE)

>

> > See below. Currently running with

>

> net.ipv4.tcp_rmem = 4096 87380 184762

> net.core.rmem_max = 184761

1) You didn't describe value of the net.core.rmem_default parameter.
Try to increase one.

2) Can you check patch from thread

<http://forum.openvz.org/index.php?t=msg&goto=9018> ?

>

> > Have you seen any suspicious messages in dmesg related to this issue?

>

> > Indeed, dmesg shows

>

> TCP: Treason uncloaked! Peer 91.4.192.15:52306/80 shrinks window 204385904:204399829.
Repaired.

> TCP: Treason uncloaked! Peer 91.4.192.15:52306/80 shrinks window 204385904:204399829.
Repaired.

> TCP: Treason uncloaked! Peer 91.4.192.15:52306/80 shrinks window 204385904:204399829.
Repaired.

> TCP: Treason uncloaked! Peer 62.158.79.38:1883/80 shrinks window

1835842591:1835855659. Repaired.
> TCP: Treason uncloaked! Peer 91.12.168.104:2918/80 shrinks window 2946649572:2946651606. Repaired.
> TCP: Treason uncloaked! Peer 217.238.252.24:49584/80 shrinks window 3063558791:3063563147. Repaired.
> TCP: Treason uncloaked! Peer 87.167.76.47:52766/80 shrinks window 867053615:867090259. Repaired.
> TCP: Treason uncloaked! Peer 91.11.116.53:2475/80 shrinks window 3699395483:3699409070. Repaired.
> TCP: Treason uncloaked! Peer 91.11.116.53:2476/80 shrinks window 3699833370:3699846438. Repaired.
> TCP: Treason uncloaked! Peer 91.11.116.53:2475/80 shrinks window 3699403648:3699409070. Repaired.
> TCP: Treason uncloaked! Peer 91.11.116.53:2476/80 shrinks window 3699841562:3699846438. Repaired.
> TCP: Treason uncloaked! Peer 79.212.215.179:2859/80 shrinks window 400801346:400802798. Repaired.
> TCP: Treason uncloaked! Peer 80.139.188.46:3646/80 shrinks window 156197829:156205177. Repaired.
> TCP: Treason uncloaked! Peer 80.139.188.46:3646/80 shrinks window 156197829:156205177. Repaired.
> TCP: Treason uncloaked! Peer 80.140.219.65:49793/80 shrinks window 4253417104:4253432263. Repaired.
> TCP: Treason uncloaked! Peer 80.140.219.65:49793/80 shrinks window 4253417104:4253432263. Repaired.
> TCP: Treason uncloaked! Peer 80.140.219.65:49793/80 shrinks window 4253417104:4253432263. Repaired.
>
> That's from the hardware node after running for three days. I'm sorry I
> cannot relate the messages to the dropout times - the messages do not
> show up in the syslogs (yet - I'm new to the server).
>
> Thanks & regards - Max.
>
>
> # cat /proc/net/netstat
>
> TcpExt: SyncookiesSent SyncookiesRecv SyncookiesFailed EmbryonicRsts PruneCalled
RcvPruned OfoPruned OutOfWindowIcmps LockDroppedIcmps ArpFilter TW TWRcycled
TWKilled PAWSPassive PAWSActive PAWSEstab DelayedACKs DelayedACKLocked
DelayedACKLost ListenOverflows ListenDrops TCPPrequeued TCPDirectCopyFromBacklog
TCPDirectCopyFromPrequeue TCPPrequeueDropped TCPHPHits TCPHPHitsToUser
TCPPureAcks TCPHPAcks TCPRenoRecovery TCPSackRecovery TCPSACKReneging
TCPFACKReorder TCPSACKReorder TCPRenoReorder TCPTSReorder TCPFullUndo
TCPPartialUndo TCPDSACKUndo TCPLossUndo TCPLoss TCPLostRetransmit
TCPRenoFailures TCPSackFailures TCPLossFailures TCPFastRetrans TCPForwardRetrans
TCPSlowStartRetrans TCPTimeouts TCPRenoRecoveryFail TCPSackRecoveryFail
TCPSchedulerFailed TCPRcvCollapsed TCPDSACKOldSent TCPDSACKOfoSent

```
TCPDSACKRecv TCPDSACKOfoRecv TCPAbortOnSyn TCPAbortOnData TCPAbortOnClose
TCPAbortOnMemory TCPAbortOnTimeout TCPAbortOnLinger TCPAbortFailed
TCPMemoryPressures
> TcpExt: 0 0 0 15666 390 7 0 3 0 0 0 0 0 0 0 350 518998 283 5722 6044 6044 13385260
860306 145918454 0 13263422 87351 9290224 12159160 1576 60817 31 631 240 92 198 367
2815 1284 1383 34266 20 369 13804 1501 90841 8390 24116 28574 309 7175 2 12934 5823 2
14577 291 0 751 5077 0 2285 0 0 0
>
> # cat /proc/net/snmp
>
> Ip: Forwarding DefaultTTL InReceives InHdrErrors InAddrErrors ForwDatagrams
InUnknownProtos InDiscards InDelivers OutRequests OutDiscards OutNoRoutes ReasmTimeout
ReasmReqds ReasmOKs ReasmFails FragOKs FragFails FragCreates
> Ip: 2 64 55086729 0 0 0 0 55079918 59160345 0 0 22 64 0 22 0 0 0
> Icmp: InMsgs InErrors InDestUnreachs InTimeExcds InParmProbs InSrcQuenchs InRedirects
InEchos InEchoReps InTimestamps InTimestampReps InAddrMasks InAddrMaskReps OutMsgs
OutErrors OutDestUnreachs OutTimeExcds OutParmProbs OutSrcQuenchs OutRedirects
OutEchos OutEchoReps OutTimestamps OutTimestampReps OutAddrMasks OutAddrMaskReps
> Icmp: 10852 99 2358 0 0 2 0 8401 0 0 0 0 8472 0 71 0 0 0 0 8401 0 0 0 0
> Tcp: RtoAlgorithm RtoMin RtoMax MaxConn ActiveOpens PassiveOpens AttemptFails
EstabResets CurrEstab InSegs OutSegs RetransSegs InErrs OutRsts
> Tcp: 0 0 0 0 487897 1752675 1326 25489 45 55038781 59121501 195520 397 7118
> Udp: InDatagrams NoPorts InErrors OutDatagrams
> Udp: 30214 71 0 30372
>
>
--
Thank,
Vitaliy Gusev
```
