

Load Generator Performance


An essential element of the controller or agent is the load generator that handles request /response processing. It needs a sufficient number of .NET threads to timely send and process HTTP messages.

The system creates the necessary number of threads automatically, depending on the test complexity, the number of VUs and the test machine's capacity. The actual number of threads allocated to StresStimulus process is controlled by the operating system, and depends on the load generator's needs to issue requests.

By default, before the test begins, StresStimulus allocates the number of .NET threads equal to the maximum number of VUs expected in this test.

During test execution, one VU may consume more than one thread to handle parallel requests. If StresStimulus needs more resources, the OS will gradually increase threads to the sufficient level. On some systems, however, OS fails to ramp-up threads fast enough. This may result in delays in issuing requests after receiving a response at some moments during the test execution. You can check if such delays exist in your system by analyzing a page or transaction waterfall chart.

If you see a gap between the moment of receiving a response and issuing a subsequent request, your system may need more threads. To address the situation, you can manually adjust the initial number of threads, which is a property of a load agent.



Note: The gap between a response and a subsequent request can be caused by different reasons. For example, StresStimulus intentionally delays subsequent request to comply with the emulated browser characteristics. Each browser type has a limited number of connections-per-host and connections-per-proxy. When the number of connection exceeds this limit, StresStimulus will pause requests to realistically mimic a physical browser behavior.

If you need to

The screenshot shows the StresStimulus configuration window. On the left is a 'Workflow Tree' with options like 'Record Test Case...', 'Build Test Case', 'Test Profile', 'Other Options', 'Load Agents', 'Monitoring', 'Result Storage', 'Pass/Fail Configuration', 'Advanced Options', 'Script Editor', 'Run and Monitor Test...', and 'Analyze Results'. The 'Load Agents' section is expanded, showing a list of agents: 'Local (weight: 1, VUs: 12)' and 'POWEREDGE (weight: 1, VUs: 13)'. The 'POWEREDGE' agent is selected, and its properties are displayed in a table below.

Agent Name	POWEREDGE
StresStimulus version	4.5.5956
Host / IP Address	10.0.2.121
Mix Weight	1
Allocate at least one thread per VU	Yes
Minimum starting threads	50

Minimum starting threads
To increase the Starting thread count, set this property to a higher number. This property will be disregarded if it is smaller than the starting thread count determined as explained in the previous property.

a
d
s
fa
st
e
r,
w
h
i
c
h
m
a
y
b
e
n
e
c
e
s
s
a
r
y
w
h
e
n
r
u
n
n
i
n
g
la
r
g
e-
s
c
al
e
te
st
s
o
n
hi
g
h
p
e
r
f
o
r
m
a
n
c
e
h
a
r
d
w
a
r
e,
s
et
th
e
**M
i
n
i
m
u
m**

s
t
a
r
t
i
n
g
t
h
r
e
a
d
s
p
r
o
p
e
r
t
y
t
o
t
h
e
l
e
v
e
l
1
2
0
%
-
1
5
0
%
o
f
t
h
e
n
u
m
b
e
r
o
f
v
i
r
t
u
a
l
u
s
e
r
s
. T
h
a
t
s
h
o
u
l
d
b
o
o
s
t
t
h
e
p
e
r
f
o
r
m
a
n
c
e
o
f

y
o
u
r
l
o
a
d
g
e
n
e
r
a
t
o
r.

A
f
t
e
r
t
h
a
t
,
r
e
r
u
n
y
o
u
r
t
e
s
t
t
o
c
h
e
c
k
i
f
s
u
c
h
a
n
i
n
c
r
e
a
s
e
w
a
s
s
u
f
f
i
c
i
e
n
t
.
Y
o
u
m
i
g
h
t
n
e
e
d
t
o
t
r
y
s
e
v

er
al
di
ff
er
ent
num
bers
until
your
load
ag
gener
ator
recei
ves
suffi
cient
resources.

This
property
is
configured
for
each

h
l
o
a
d
a
g
e
n
t
i
n
d
i
v
i
d
u
a
l
l
y
.
N
o
t
e
t
h
a
t
s
e
t
t
i
n
g
a
t
o
o
h
i
g
h
n
u
m
b
e
r
o
f
t
h
r
e
a
d
s
m
a
y
l
e
a
v
e
i
n
s
u
f
f
i
c
i
e
n
t
r
e
s
o
u
r
c
e
s
f
o
r
t
h
e
O
S
a
n
d
c
a
n

slow
down
your
com-
puter.

I n f o: One VU is not equal to one thread. Threads available to the process (thread pool) c

a
n
b
e
u
s
e
d
b
y
di
ff
e
r
e
nt
V
U
s
at
di
ff
e
r
e
nt
ti
m
e
s.
E
a
c
h
V
U
c
a
n
is
s
u
e
m
ul
ti
pl
e
c
o
n
c
u
rr
e
nt
r
e
q
u
e
st
s
(c
o
n
n
e
ct
io

n
s)
,
d
e
p
e
n
d
i
n
g
o
n
y
o
u
r
t
e
s
t
c
a
s
e
a
n
d
t
h
e
c
o
n
n
e
c
t
i
o
n
l
i
m
i
t
o
f
t
h
e
s
e
l
e
c
t
e
d
b
r
o
w
s
e
r.
F
o
r
e
x
a
m
p
l
e,

I
E
9
c
a
n
o
p
e
n
u
p
to
6
c
o
n
n
e
ct
io
n
s.

N
o
t
e:
In
the
p
r
e
vi
o
u
s
v
e
rs
io
n
s,
the
r
e
a
d
al
lo
c
at
io
n
al
g
o
r
i
t
h
m
r
e
l
ie
d
o
n
.
N
E

The screenshot shows a software interface with a 'Main' window. On the left is a 'Workflow Tree' with the following items: Record Test Case..., Build Test Case, Test Profile, Other Options, Load Agents, Monitoring, Result Storage, Pass/Fail Configuration, Advanced Options, Script Editor, Run and Monitor Test..., and Analyze Results. The 'Load Agents' item is selected. The main area displays configuration for 'Local (weight 1, VUs: 1)'. The configuration table is as follows:

Mix Weight	1
Allocate at least one thread per VU	Yes
Minimum starting threads	1
VUs	1

Below the table, there is a section titled 'Allocate at least one thread per VU' with the following text: 'Starting thread count is the number of threads created automatically when the test is launched. By default, it is equal to the number of VUs. If more threads are needed, the load engine will gradually create more threads while checking available system resources. Select No to allow .NET Framework to allocate a default (typically smaller) number of starting threads.'

T
F
r
a
m
e
w
o
r
k
's
d
e
f
a
u
l
t
a
l
g
o
r
i
t
h
m
o
f
a
l
l
o
c
a
t
i
n
g
t
h
r
e
a
d
s.
A
s
a
r
e
s
u
l
t,
t
h
e
s
t
a
r
t
i
n
g
t
h
r
e
a
d
c
o
u
n
t
w
a
s
t
y
p
i
c
a
l
l
y

s
m
a
l
l
e
r
t
h
a
n
t
h
e
n
u
m
b
e
r
o
f
v
i
r
t
u
a
l
u
s
e
r
s
. F
o
r
b
a
c
k
w
a
r
d
c
o
m
p
a
t
i
b
i
l
i
t
y
,
t
h
e
o
p
t
i
o
n
t
o
u
s
e
t
h
e
l
e
g
a
c
y
t
h
e
r
e
a
d
a
l
l
o
c
a
t
i
o

n al g o r i t h m e x i s t . T o u s e i t , s e t p r o p e r t y A l l o c a t e a s t o n e t h r e a d p e r V U t o N o .

T y p i c a l l y t h e l o a d e n g i n e i s o p t i m

Workflow Tree

Graph sample rate (s)	1
Pre-run Command Line	
Pre-run Command Timeout	10
MIME Types requested sequence	
Enable Dynatrace integration	No
Optimize load engine for	Time Resolution

Optimize load engine for
Choose Performance to handle more virtual users on a single load generator. Choose Time Resolution to measure response time under 0.1 second with more precision.

ized for performance to handle them more virtually users on a single load generator at a time. However, if your test has many users

rc
e
r
e
q
u
e
s
t
s
t
h
a
t
h
a
v
e
a
s
m
a
l
l
e
x
p
e
c
t
e
d
r
e
s
p
o
n
s
e
t
i
m
e
w
a
n
t
t
o
c
h
a
n
g
e
t
h
e
l
o
a
d
e
n
g
i
n
e
o
p
t
i
m
i
z
a
t
i
o
n
f
o
r
"
T
i
m
e
R
e
s
o
l
u
t
i
o
n

". This will be a blameless sunrise response under 0.1 second with more precision, however, will have a performance

n
c
e
i
m
p
a
c
t
w
i
t
h
m
o
r
e
v
i
r
t
u
a
l
u
s
e
r
s.