

Remote Analytics Agent Sizing

On this page:

- [Analytics Agent Sizing Based on Event Volume](#)

Related pages:

- [Deployment Options and Scenarios](#)
- [Installing Agent-Side Components](#)
- [Event Service Sizing and Capacity](#)

★

★

★

★

★

★

This page contains information to help you estimate the hardware requirements for a remote Analytics Agent deployment where a single Analytics Agent is aggregating transaction events from multiple APM Application Agents. This guide covers transaction analytics only. There are two basic issues regarding sending analytics data to a remote Analytics Agent:

- How many APM agents can report to one remote Analytics Agent?
- What are the machine requirements for hosting the remote Analytics Agent?



Don't extract the numbers in this guide to sizing for Log Analytics because the Analytics Agent must be installed on the local machine to capture Log Analytics.

Analytics Agent Sizing Based on Event Volume

Based on our testing, the volume of events being sent to the Analytics Agent is the limiting factor in determining how many APM agents can report to one remote Analytics Agent.

The tests were conducted on virtual hardware and programmatically generated work load. Real world work loads may vary. To best estimate your hardware sizing requirements, carefully consider the traffic patterns in your application and test in a test environment that closely resembles your production application and user activity.

Calculating Analytics Event Volume

One business transaction can traverse many tiers. In each tier, one business transaction traverses one node. One node produces one request per business transaction when the transaction is synchronous. For async transactions, multiple events may be generated by a node for a single request. One request equals one analytics event. To calculate how many events a business transaction generates, you need to count the number of tiers/nodes that are sending data into the Analytics Agent.

You can estimate the number of events using the following formula:

One business transaction generates events at a rate = **calls per minute** times the **number of tiers** reporting analytics data for the business transaction.

In simple terms: #events for one business transaction = calls per minute times # of tiers.

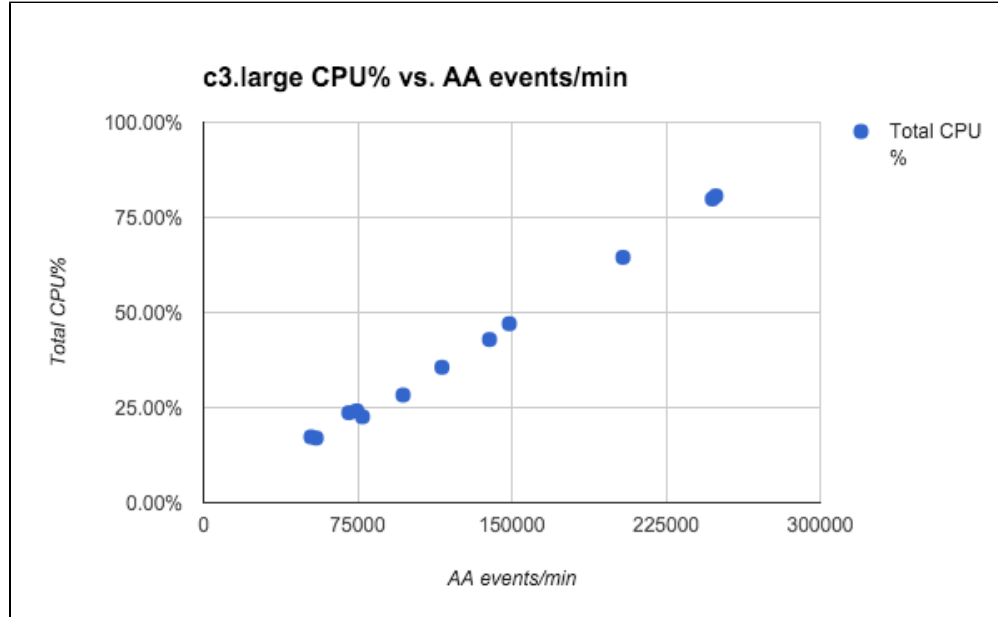
Characteristics of the Amazon EC2 Instance Types

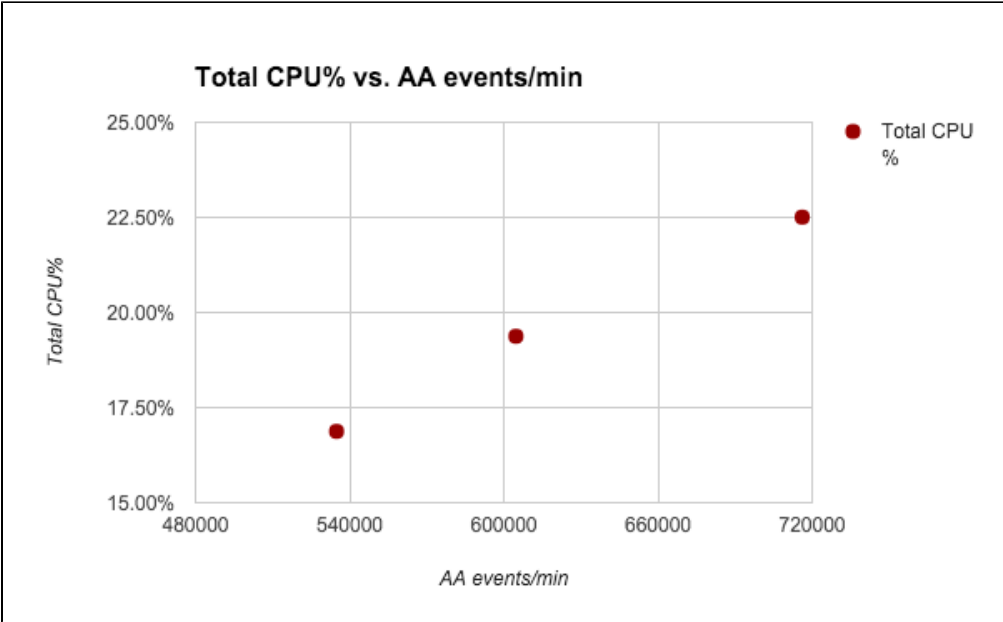

For complete information on Amazon EC2 instance types, see <https://aws.amazon.com/ec2/instance-types/>.

The testing was performed using c3.large, c3.xlarge, and c4.4xlarge.

Model	vCPU	Mem (GiB)	SSD Storage (GB)
c3.large	2	3.75	2 x 16
c3.xlarge	4	7.5	2 x 40
c4.4xlarge	16	30	EBS-Only

Testing Results





Raw Data

Analytics Agent Host Machine	Analytics Agent events/min	Total CPU%	JVM Heap (Mb)
c3.large	52313	17%	468
c3.large	54831	17%	470
c3.large	70746	24%	475
c3.large	74541	24%	477
c3.large	77344	23%	487
c3.large	97074	28%	512
c3.large	115999	36%	519
c3.large	139143	43%	526

c3.large	148782	47%	587
c3.large	204073	65%	527
c3.large	247543	80%	624
c3.large	249261	81%	637
c3.xlarge	196288	33%	518
c3.xlarge	254586	44%	727
c3.xlarge	302689	51%	497
c3.xlarge	336879	58%	913
c3.xlarge	372515	65%	1024
c3.xlarge	513598	91%	922
c3.xlarge	478954	78%	922
c3.xlarge	420000	69%	979
c3.xlarge	376034	61%	1024
c3.xlarge	318000	52%	1024
c3.xlarge	258000	43%	1024
c3.xlarge	198000	32%	1024
c3.xlarge	144000	22%	1024
c4.4xlarge	534900	17%	552
c4.4xlarge	604725	19%	841
c4.4xlarge	716141	23%	1024