SIZING ESTIMATE

John Rupf

Please send feedback to cs-mars@cisco.com

INTRODUCTION

This document introduces parameters to consider when sizing the CS-MARS model to the customer's requirements

TABLE OF CONTENTS

INTRODUCTION

A customer has some number of firewalls, servers, routers, and IPS devices and would like to estimate the approximate size for CS-MARS. Specifically how many and what model of CS-MARS are required.

The analysis of the sizing problem may be broken up into four parts: the number of sites, the requirement for high availability, the events per second, and the online storage space needed.

If the customer's devices are distributed between sites and the cost of transferring data between sites is a consideration, then two or more smaller CS-MARS appliances may be appropriate for the job. If all of the customer's devices are at one site, or data transfer between sites is cheap, a single larger appliance may be the right selection.

If the customer has a requirement for high availability, then the RAID level of each model may eliminate some models from consideration.

If the devices are generating a great deal of traffic to CS-MARS then a CS-MARS with the capacity to handle more events per second is required. If the number of devices is expected to grow, then a larger CS-MARS may be appropriate.

Depending on the amount of storage space the customer needs to keep data in online storage for a given period of time, some models may be eliminated because the model does not have the required online storage capacity.

Answering the following questions will help you size the CS-MARS requirements.

QUESTIONS

1. How many sites will CS-MARS support?
2. Does CS-MARS need to be highly available?
3. How many events will CS-MARS receive per second?
4. How many months of data must CS-MARS make available online?

Sites

Devices may be all in one site or distributed between sites. If the devices are in one site or the cost of transferring data between sites is low, then one CS-MARS appliance may be able to handle the job.

Sites may belong to separate administrative units where smaller boxes that each administrative unit could own may be a better solution than one larger, shared CS-MARS.

Do not distribute devices to CS-MARS by device type. CS-MARS must be able to follow the packet through the network.

High Availability

If the customer has a requirement for high availability, then begin at RAID 1 + 0 or a model 100. The models having RAID 1 + 0 also have dual redundant power supplies.

The 110 and 210 have cache battery backup as well.

Model	RAID Level	Power Supply	Battery Backup
25R	None		
25	None		
55	0		
110R	1 + 0	Redundant	Built-in
110	1 + 0	Redundant	Built-in
210	1 + 0	Redundant	Built-in

Table 1 High Availability

If high availability is a requirement you should recommend at least the 110R.

Events per Second

System Log Messages

Begin with the number of events per day. If the traffic is uniform throughout the day, then convert the events per day directly into events per second.
Calculate estimated number of events per second.

Events/hour = Events/day * 1 day / 24 hours
Events/second = Events/hour * 1 hour/3600 second

For example, assume 5 000 000 Events/day
Events/second ~= 58

If the traffic on the network is not uniform throughout the day, then try to estimate the events per second for the busiest part of the day. For example, assume 80% of the events occur over a 10 hour period during the day.
Calculate the estimated number of events per second.

Events/hour = Events / day * 0.8 / 10 hours
Events/second = Events / hour * 1 hour/3600 second

For example, assume 5 000 000 Events/day
Events/hour = 5 000 000 * 0.8 / 10 = 400 000
Events/second ~= 111

SDEE

Use IME statistics to estimate the rate of SDEE events per second.

Composite EPS

Calculate a composite EPS value by adding 3 times the SDEE EPS to the system log message EPS.
For example, to calculate the estimated EPS load on the appliance where the system log EPS is 250 and the SDEE EPS is 100 then
Composite EPS = 250 + 3 * 100 = 550 EPS

A CS-MARS 55 at 900 Events/second sustained is the first appliance in table 2 that can sustain 550 EPS. To allow room for traffic growth and bearing in mind that during an attack, traffic could go up by a factor of 5 or more you should recommend the 110R.
Allow more room if CS-MARS will run custom reports.

Peak Sustained Peak Sustained Peak Sustained

	Syslog or NetFow-Store EPS	Syslog or NetFow-Store EPS	NetFlow non-store EPS (V5)	NetFlow non-store EPS (V5)	ASA NF-non-store EPS (V9)	ASA NF-non-store EPS (V9)	IPS SDEE per sec	IPS SDEE per sec
Platform	Peak	Sustained	Peak	Sustained	Peak	Sustained	Peak	Sustained
MARS 25R	75	45	1 500	300	750	450	27	16
MARS 25	750	450	15 000	2 000	7 500	3 000	270	162
MARS 55	1500	900	30 000	6 000	15 000	9 000	540	324
MARS 110R	4 500	2 700	75 000	15 000	35 000	22 500	1 620	972
MARS 110	7 500	4 500	150 000	30 000	75 000	45 000	2 700	1 620
MARS 210	15 000	9 000	300 000	60 000	150 000	90 000	4 000	2 400

Table 2 Events and Netflow by Device

Online Storage Space

CS-MARS can use two types of storage: online and archival. The online storage is available to the database. The archival storage contains the database archives. For sizing we need only consider the online storage space as a function of the model.

After allowing for overhead the actual storage space available for storing events will be smaller than the nominal storage space of the hard drive.

Model	Maximum Events	Maximum Days at 100 EPS
25	1.3 * 10E8	~ 15
55	2.7 * 10E8	~ 31
110	2.4 * 10E9	~ 270
210	3.6 * 10E9	~ 410

Table 3 Useable space

In table 3 we see the total disk space for each model along with the estimated available storage space. The Estimated Usable Storage Space column shows the dedicated event space that is available after deduction of space for overhead.

At this point we need an estimated number of byes per week in order to continue with the sizing.

For this example assume 7 GB of events per week and that our customer wants to have four weeks of online data.

GB = week * GB/week

For example, at 7 GB per week
GB = 4 * 7 GB = 28 GB
To keep 4 weeks of data would require 28 GB

Looking at table 3 we can see that a CS-MARS 20 would be able to handle these storage requirements and have room to handle growth or the increased volume due to an attack.

Offline Storage

6 GB = 10 EPS * 1 year * 200 bytes/ message * 1/10 compression, then

Assuming 3 months is 1/4 of a year, we estimate about 1.5 GB for 3 months at 10 EPS, 0.15 GB at 1 EPS.
We need an estimate of the EPS. We typically recommend setting up a Kiwi syslog server if there is no syslog server currently available.
Multiply the EPS times 0.15 GB to get the 3 month estimate.
This is not a science. This is an estimate. During an attack the EPS may go much higher. This assumes syslog messages of 200 bytes in length and compressed to 1:10. This will vary by message. Traffic is usually not the same

throughout the day or the year. If the traffic on the network increases over time due to growth or the addition of more reporting devices, this estimate will be low.

OBTAINING EVENTS PER SECOND AND BTYES OF EVENTS PER WEEK

The most accurate way to determine the number of events per second needed is to put a system log daemon like the Kiwi syslog daemon in the customer's network at the IP address where they propose putting the CS-MARS.
This product can be downloaded from http://www.kiwisyslog.com/index.php.
Get the number of bytes and events for one day as an estimate.

ESTIMATION RULES OF THUMB

This section is about making guesses based on the device.

Estimating Bytes of Events per Week from Events per Day

Assume that the customer has an estimate of the number of events per day, but not the bytes per week. First, we must estimate the size of each event. If the events are primarily from routers and firewalls, then an estimated event size of 200 bytes is about right. If there are a number of IPS/IDs devices then an estimated event size of 500 bytes is probably better. Keep in mind that CS-MARS will store up to 1.5 MB including any captured packets.

Device Type	Estimated Event size
Firewall	200 bytes
IPS/IDS	500 bytes

Table 4 Estimated Event Size

For this example assume 200 bytes per event, 5 000 000 events per day, and one month of online data. This data is strictly for purposes of illustration.

Bytes/day = Events/day * 200 bytes/Event
Bytes/week = Bytes/day * 7 days/week

For example, 5 000 000 Events/day at
Bytes/day = 5 000 000 * 200 = 1 GB
To keep 4 weeks of data would require 28 GB

Any connection to the Internet is constantly being tried for weaknesses. The customer's edge routers and firewalls will turn back much of this probing so the farther these devices are from the edge the fewer extraneous events they will produce. These events typically result in false positives anyway as they are dropped at the edge. They can be voluminous. The devices in the campus will give a lower volume of more relevant results. We can tune the MARS to drop false positives and tune the firewalls to reduce production of false positives.

Failover firewalls may be run in active/active or active/standby mode. Count as one fire in active/standby mode.

Device Class	EPS Estimate	Estimated Bytes / event	Total bytes per week
Windows Server	Limited to 100 when MARS polls	200	Customer supplied estimate

Windows Domain Controller	Limited to 100 when MARS polls	200	Customer supplied estimate
Firewall	2 EPS for each connection per second	200	Customer supplied estimate
IPS/IDS	Customer supplied estimate	500	Customer supplied estimate
Router	Customer supplied estimate	200	Customer supplied estimate

Table 5 Guesses by Device Class

USING THE SIZING CALCULATOR

Raymond Jett has had good results using the online sizing calculator by using the following three rules.

1. Use a system log server to get the number of events for one day. Normalize to the events per day to events per second.
2. Give the actual usage of an interface rather than its nominal value. For example, if a PIX 501 is connected to a T1 use 1.54 Mbps, rather than 10/100 Mbps.
3. Move up one model from the model that the calculator recommends. If the calculator recommends a CS-MARS 100e, then move up to a CS-MARS 100.

REFERENCES

Security Threat Mitigation and Response: Understanding Cisco Security MARS by Dale Tesch, Greg Abelar

Publisher: Cisco Press Pub Date: September 28, 2006
Print ISBN-10: 1-58705-260-1
Print ISBN-13: 978-1-58705-260-6

Security Monitoring with Cisco Security MARS by Gary Halleen, Greg Kellogg

Publisher: Cisco Press Pub Date: July 6, 2007
Print ISBN-10: 1-58705-270-9
Print ISBN-13: 978-1-58705-270-5