

Distributing SAP Load using BIG-IP Advanced Monitoring



Nojan Moshiri, 2011-09-03



Several recent forum posts on [DevCentral](#) forums have commented on the fact that

SAP Landscapes often have asynchronous batch jobs that cause higher CPU loads on certain servers. This causes problems for application delivery controllers because load balancing methods are typically based on connection counts. Picture the scenario where one connection causes a big CPU or memory spike and then goes away. Now you have the same number of new connections coming into the server while one is slammed.

The solution to this problem is relatively straightforward and I recently documented this for everyone in our “Deploying F5 Networks with SAP NetWeaver” deployment guide, located here: [SAP NetWeaver and Enterprise SOA: Enterprise Portal \(BIG-IP v10.1, WOM, Edge, WA\)](#). The solution is based around using SNMP in conjunction with application based monitors. The BIG-IP SNMP monitor provides the ability to perform dynamic load balancing based on CPU, memory or disk utilization while the advanced monitors test the J2EE stack, the authentication system and the database. With this combination, SAP administrators should be able to sleep better at night knowing that their customers and users are getting to a live system that best prepared to service the request.

So, how does layer monitoring work? If you are not aware, it’s possible to have two monitors for a particular pool or node. In the UI, it looks like this:

The screenshot shows the configuration page for a node named 'SAP-Discovery' with IP address 10.133.63.50. The 'General Properties' section shows the node is available and has a health monitor named 'SAP-CPU'. The 'Configuration' section shows the health monitor is set to 'Node Specific' and 'Active'. The 'Select Monitors' section shows two monitors: 'SAP-CPU' and 'icmp'. The 'Availability Requirement' is set to 'All Health Monitor(s)', the 'Ratio' is 1, and the 'Connection Limit' is 0. Buttons for 'Update' and 'Delete' are visible at the bottom.

General Properties	
Address	10.133.63.50
Name	SAP-Discovery
Partition	Common
Availability	Available (Enabled) - Node address is available
Health Monitors	SAP-CPU
Current Connections	0
State	<input checked="" type="radio"/> Enabled (All traffic allowed) <input type="radio"/> Disabled (Only persistent or active connections allowed) <input type="radio"/> Forced Offline (Only active connections allowed)

Configuration													
Health Monitors	Node Specific												
Select Monitors	<table border="1"><thead><tr><th>Active</th><th>Available</th></tr></thead><tbody><tr><td>SAP-CPU</td><td>gateway_icmp</td></tr><tr><td>icmp</td><td>https_443</td></tr><tr><td></td><td>real_server</td></tr><tr><td></td><td>snmp_dca</td></tr><tr><td></td><td>tcp_echo</td></tr></tbody></table>	Active	Available	SAP-CPU	gateway_icmp	icmp	https_443		real_server		snmp_dca		tcp_echo
Active	Available												
SAP-CPU	gateway_icmp												
icmp	https_443												
	real_server												
	snmp_dca												
	tcp_echo												
Availability Requirement	All Health Monitor(s)												
Ratio	1												
Connection Limit	0												

Update Delete

In this example there are two monitors, SAP-CPU and ICMP. In the real world, ICMP would be replaced with the advanced application monitor. So, what does the SNMP monitor configuration look like:

Local Traffic » Monitors » SAP-CPU

Properties
 Instances

General Properties

Name	SAP-CPU
Partition	Common
Type	SNMP DCA

Configuration: **Advanced**

Interval	<input type="text" value="10"/> seconds				
Time Until Up	<input type="text" value="0"/> seconds				
Timeout	<input type="text" value="30"/> seconds				
Community	<input type="text" value="public"/>				
Version	<input type="text" value="v2c"/>				
Agent Type	<input type="text" value="WIN2000"/>				
CPU Coefficient	<input type="text" value="1.5"/>				
CPU Threshold	<input type="text" value="80"/> %				
Memory Coefficient	<input type="text" value="1.0"/>				
Memory Threshold	<input type="text" value="0"/> %				
Disk Coefficient	<input type="text" value="2.0"/>				
Disk Threshold	<input type="text" value="10"/> %				
Variables	<table border="1"> <thead> <tr> <th>Name</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td><input type="text"/></td> <td><input type="text"/></td> </tr> </tbody> </table> <p>Add</p>	Name	Value	<input type="text"/>	<input type="text"/>
Name	Value				
<input type="text"/>	<input type="text"/>				

Here we have an SNMP setup that is set at a CPU Threshold of 80%, a memory Threshold of 0% and a Disk Threshold of 10%. Obviously this is from my testing to insure the monitor is working properly. What this defines is that if the disk is more than 10% full, or the memory is being utilized at 0% or the CPU is being utilized at over 80%, then de-weight the amount of new connections that get sent to this node(server). The coefficients allow further granular control over the traffic weighting determination. This is not a config you would probably run in production, but it's great for testing!

By logging into the BIG-IP advanced shell and enabling logging, I can see exactly what weight is being assigned. This is accomplished through the command:

```
bigpipe db Snmp.SNMPDCA.Log true and then by tailing the snmpdca.log located in /var/tmp :
```

```
tail -f /var/tmp/snmpdca.log
```

There you have it. Now all we have to do is change the load balancing mechanism for the pool to be based on dynamic, apply the advanced application monitor, and we have a fully dynamic decision making system. You can play with the Thresholds and Coefficients until you have a desired mix. The SNMP monitor will not mark a host down, but it will set the weight (between 1 and 100) in a manner that very few connections will get to a node that has exceeded all tresholds.

A quick note on the advanced health monitor. I can't stress how important it is to have layered monitoring in this and other dynamic load balancing scenarios. Especially in an SAP NetWeaver J2EE stack installation (or even a dual stack implementation) many things can go wrong. Just because the CPU, memory and disk are normal, doesn't mean that your J2EE stack hasn't crashed, or that your authentication system has gone down. By layering monitors, you cover all BASIS. :-)

I hope this post has been helpful, and as always, please email me if you have any questions. Remember that detailed installation instructions including step-by-step configuration is in the deployment guide linked at the top, or through [f5.com](#) ---> Resources -- > Deployment Guides ---> SAP NetWeaver and Enterprise SOA: Enterprise Portal (BIG-IP

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com