

Results for: OpenStak/LBaaS - Project-user use cases

In addition to selecting one of the option, an optional text field is provided to provide more information.
Please answer to each of the use cases using the following terms:

- **Must have** - This use case applies strongly to our organization and we will not use the product if we can't support this use case.
- **Want** - We do not presently support this use case but would like to. Our users are asking for it.
- **Don't care** - We don't support this use case and don't think our users want it. But we don't care if the product supports this use case.
- **"DO NOT WANT"** - We will not support this use case and will not use this product if it does. It is advised to provide further information in the accompanying text field

1) Please identify yourself

Name	Organization
Basil Baby	Comcast
Stephen Balukoff	Blue Box Group
Fima Lifshitz	Radware/Professional Services
Trevor Vardeman	Rackspace
Carlos D. Garza	Rackspace
Kevin Fox	PNNL
Youcef Laribi	Citrix
Dustin Lundquist	Blue Box Group
Tim Cuddy	HP
Jorge Miramontes	Rackspace
Jay Pipes	Mirantis
Samuel Bercovici	Radware

2) A project-user wants to make his web based application highly available. He has n VMs deployed on the same private subnet/network. Each VM is installed with a web server (ex: apache) and content. The project-user wishes to represent them to the application users as a web application available via a single IP.

	Response (%)	Responses
Must have	91.67	11
Want	8.33	1
Don't care	0.00	0
"DO NOT WANT"	0.00	0
	Answered Question	12
	Skipped Question	0

Additional Information (for the selected choice):

vip -> listener (http) -> pool

This is the essence of lbaas :)

3) A project-user web application has two parts. the first part which contains n VMs deployed on the same private network and are available via HTTP and the second part which contains m VMs deployed on the same private network and are available via HTTPS. He wants to make application highly available and exposed to the application users via a single IP.

	Response (%)	Responses
Must have	58.33	7
Want	41.67	5
Don't care	0.00	0
"DO NOT WANT"	0.00	0
	Answered Question	12
	Skipped Question	0

Additional Information (for the selected choice):

-> listener(http) -> pool1 vip(vip-subnet, pools-subnet) -> listener(https) -> pool2



We would most likely just need to solve the case where HTTP/HTTPS traffic is being served on the same nodes. However, I'm not against this.

This isn't entirely clear where the user wants to present two ports (HTTP and HTTPS) each backed by a different pool of servers, or if wants point part of

a HTTP application to another pool of applications servers based on some layer 7 feature. Since the first is more likely I'm going with that.

The important part here is the ability to: - Have multiple lb configs for the same subnets - Support this when the VIP IP address is the same or different.

4) A project-user web application has two parts. the first part which contains n VMs deployed on the same private network and are available via HTTP and the second part which contains m VMs deployed another private network and are available via HTTPS. He wants to make application highly available and exposed to the application users via a single IP.

		Response (%)	Responses
Must have		41.67	5
Want		58.33	7
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0



Additional Information (for the selected choice):

-> listener(http) -> pool1 vip(vip-subnet, pool1-subnet, pool2-subnet) -> listener(https) -> pool2

This is the same use case as #3.

The fact that the VIP IP address is the same shouldn't mean the subnet of the VMs of the corresponding LB configs has to be the same.

5) A project-user wants to make his secured (HTTPS) web application, also available to users who tried to access the application via an unsecured channel (HTTP) by redirecting them back to the secured application. The load balancer will detect access to HTTP and will redirect to HTTPS The application members can handle the redirect logic

		Response (%)	Responses
Must have		58.33	7
Want		41.67	5
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):



-> listener(http) -> I7Policy (http redirect everything to port 443) vip(vip-subnet, pool2-subnet) -> listener(https) -> pool2

I prefer that the load balancer handle the redirect logic. I agree with #1, not necessarily #2

This seems to be two separate ways of implementing the use case. Option 1 would be preferred, but option 2 would be acceptable.

This is a nice to have feature, and can be part of the L7 processing capability of LBaaS.

6) A project-user wants to make his web based application highly available. He has n VMs deployed on the same private subnet/network. Each VM is installed with a web server (ex: apache) and content. The application requires that a transaction which has started on a specific VM will continue to run against the same VM. The application is also available to end-users via smart phones, a case in which the end user IP might change. The project-user wishes to represent them to the application users as a web application available via a single IP.

		Response (%)	Responses
Must have		66.67	8
Want		33.33	4
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):



stickiness policy using cookie persistence (cookie maybe inserted by the load balancer)

I presume we are talking about sticky sessions here, and not real transactions...

I call this session persistence.

If we are talking about Cookie-based persistence, then yes we must have it. Note that the cookie in this scenario should be managed by the application (an app cookie), as it's capable that it is the same user (via authentication) when he/she moves to a smartphone device with a different client IP.

7) A project-user wants to make his secured web based application (HTTPS) highly available. He has n VMs deployed on the same private subnet/network. Each VM is installed with a web server (ex: apache) and content. The application requires that a transaction which has started on a specific VM will continue to run against the same VM. The application is also available to end-users via smart phones, a case in which the end user IP might change. The project-user wishes to represent them to the application users as a web application available via a single IP.

		Response (%)	Responses
Must have		58.33	7
Want		41.67	5
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

either: persistence based on SSL session id or do SSL termination and persistence based on cookie

I presume we are talking about sticky sessions here, and not real transactions...

The difference with use case is the use of HTTPS. To support this use case, SSL must be offloaded to the loadbalancer.

8) A project-user has several different application parts developed by different teams and requiring different VMs. Nonetheless, all the different parts must be presented as a single secured application (2K SSL keys). The different parts will be available from "one" IP and different different URIs (ex: https://acme.com/maketplace/us, https://acme.com/maketplace/eu, and https://acme.com/checkout). Moreover, the communication between the load balancer and the checkout part must remain secured but the security can use weaker private keys (1K SSL keys).

		Response (%)	Responses
Must have		16.67	2
Want		66.67	8
Don't care		16.67	2
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0




Additional Information (for the selected choice):

SSL termination with back end re-encryption + L7 content switching

Seems very useful, but not necessary at this point.

This use case combines L7 content switching, as well as decryption-reencryption at the loadbalancer. We can probably tease these 2 aspects apart.

9) A project-user wants to make his web based application highly available. He has n VMs deployed on the same private subnet/network. Each VM is installed with a web server (ex: apache) and content. The project-user wishes to represent them as a standard web application via a single IP available on a VIP private network accessible to enterprise application users connecting to the cloud via a VPN. The project-user wishes to represent the same VMs as a secured web application (HTTPS) to users that access the application via the public internet.

		Response (%)	Responses
Must have		36.36	4
Want		36.36	4
Don't care		27.27	3
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

add notion of I3 selection rule on VIP that based on source IPs will steer traffic to an appropriate listener

This is an important use case. The same pool of VMs should be allowed to be accessed from more than one VIP.

10) A project-user wants to insert Client Certificate Information into HTTP Headers when the real-servers require information from the Certificate to verify or track the client.

		Response (%)	Responses
Must have		0.00	0

Want		66.67	8
Don't care		33.33	4
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

This is probably in case of SSL termination. Does this feature becomes part of the SSL policy, L7 content modification or property on listener.

Client certificate authentication is nice to have, but not a must.

11) A project-user wants to expose the real connecting client IP to her web based application behind a load balancing service. She can use an X-Forwarded-For header inserted by the load balancer to do this, but needs to be able to know the IP addresses from which the load balancer will originate requests to her application servers, so she can add them to a "trusted proxies" list (so she doesn't accidentally use this header when the source of the originating requests is untrusted).

		Response (%)	Responses
Must have		75.00	9
Want		16.67	2
Don't care		0.00	0
"DO NOT WANT"		8.33	1
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

This is probably in case of SSL termination. It should be a feature available as L7 content modification or a property on listener

Internal IP addresses in a cloud environment are dynamic and cloud applications should be designed to not rely on static/fixed IP addresses. For example, an application can simply check that the originating IP address is on the same subnet as the VM.

12) A project-user is paying for a private cloud in which n load balancing appliances have been deployed. She has n production load balancing services she wants to deploy, and needs a way to guarantee these don't end up on the same physical devices, and doesn't want to have to get the operator involved in this deployment decision.

		Response (%)	Responses
Must have		50.00	6
Want		33.33	4
Don't care		0.00	0
"DO NOT WANT"		16.67	2
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

this should be part of scheduling hints also with addition of flavor, etc. should review <https://wiki.openstack.org/wiki/InstanceGroupApiExtension> as an example

this is trying to bridge the operator world with the user world, and doesn't make a whole lot of sense in a cloud environment. This is managed hosting, nothing more, IMO. a "cloudy" way of doing this is to deploy a VM -- and use the regular scheduler hints for placement logic -- and have the load balancer deployed in the VM.

We currently call this cluster affinity. However, a less implementation specific term should be created that is easy to understand.

A project user is consuming services not physical devices. A cloud operator can however expose a service level which guarantees "a dedicated device", but the operator is definitely in control of what is offered.

13) A project-user is paying for a private cloud in which n load balancing appliances have been deployed. She has n-1 production load balancing services she wants to deploy on separate physical appliances, and several QA, staging and test services she wants to deploy, which should all share the same physical host (and shouldn't share hosts with any of the production services). As she may decide to alter the arrangement of the distribution of these load balancing services at any time, she doesn't want the operator to have to be involved in this deployment logic (and neither does the operator).

		Response (%)	Responses
Must have		8.33	1
Want		33.33	4
Don't care		50.00	6
"DO NOT WANT"		8.33	1
		Answered Question	12

Additional Information (for the selected choice):

I thing that tenant will not manage scheduling in such specificity. I do think the the spirit of the use case beads to be met. Scheduling hints is more appropriate and should allow better global placement and efficiency

this is trying to bridge the operator world with the user world, and doesn't make a whole lot of sense in a cloud environment. This is managed hosting, nothing more, IMO. a "cloudy" way of doing this is to deploy a VM -- and use the regular scheduler hints for placement logic -- and have the load balancer deployed in the VM.

I understand wanting to keep virtual devices on different hardware. However, why would a user care if some should be forced to be on the same hardware? By default this MAY happen, but if it doesn't what is the big deal?

I see this as a very advanced case whose need is not proven. The operator could define a service assigning N dedicated devices to a project, giving freedom to the project on how they partition the consumption of this capacity. This is analogous to sub-letting, with a recursive operator model.

14) A project-user has an HTTPS application in which some of the back-end servers serving this application are in the same subnet, and others are across the internet, accessible via VPN. He wants this HTTPS application to be available to web clients via a single IP address.

		Response (%)	Responses
Must have		0.00	0
Want		58.33	7
Don't care		41.67	5
"DO NOT WANT"		0.00	0
Answered Question			12
Skipped Question			0

Additional Information (for the selected choice):

consider adding static route(s) property to vip in addition to vip-subnet and pool(s)-subnet

Edge case.

This could be a useful feature for "cloud bursting" where some of the backends of an application are on-premise and others are in an OpenStack cloud.

15) A project-user has an HTTP-based application available through a load balanced service on a single IPv4 address. He wants to make the same application available via an IPv6 address, too.

		Response (%)	Responses
Must have		66.67	8
Want		16.67	2
Don't care		16.67	2
"DO NOT WANT"		0.00	0
Answered Question			12
Skipped Question			0

Additional Information (for the selected choice):

vip should have either or both ipv4 and ipv6 address

Absolute must have.

Yes, multiple VIPs for the same app is an important use case, whether it is for using different protocols (HTTP versus HTTPS), different IP addresses (internal vs external), or different IP versions (IPv4 vs. IPv6)

16) A project-user has an HTTP-based application available through a load balanced service. He would like to designate some of his back-end application servers as "backup servers" to be accessed only when the primary application servers are unavailable.




		Response (%)	Responses
Must have		41.67	5
Want		41.67	5
Don't care		16.67	2
"DO NOT WANT"		0.00	0
Answered Question			12
Skipped Question			0

Additional Information (for the selected choice):

I call this ACTIVE/PASSIVE nodes or PRIMARY/SECONDARY nodes. This is within a pool and not between pools.

Modern ADC appliances support this feature, and it would be a nice feature to have for users.



17) A project-user has an HTTP-based application available through a load balanced service. He would like a customized "error 503" page displayed whenever all the back-end servers are unavailable (including any backup servers).

		Response (%)	Responses
Must have		41.67	5
Want		25.00	3
Don't care		33.33	4
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

Being able to return a meaningful error page to a user rather a "connection timeout" would be very helpful for a better user experience.

18) A project-user has a web-based application in which they'd like the load balancer to speak HTTPS with the web clients, but HTTP with the back-end servers. Back-end servers should be made aware via the X-Proto header whether connections between the web clients and the load balancer service were done via HTTPS.

		Response (%)	Responses
Must have		50.00	6
Want		50.00	6
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

this is a hugely popular option for deploying web applications, both for performance and for simplicity reasons.

This is SSL termination.

Nice to have. Generally part of L7 content processing.

19) A project-user has an HTTP-based application behind a load balancing service, but needs to make sure no more than n clients connect with each back-end server at a time to ensure they don't get overloaded.




		Response (%)	Responses
Must have		8.33	1
Want		83.33	10
Don't care		8.33	1
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

I call this connection throttling.

Being able to limit the number of connections that a VM can handle is nice to have, and modern ADCs support this feature. A twist on this use case, is the ability to define VMs or backend servers that are only used for overflow of capacity (spare capacity).

20) A project-user has HTTP-based application behind a load balancing service. She needs to make sure that only back-end service that pass a HTTP-based health check are accessible to web clients. However, this health-check is only accessible via a URL that requires HTTP Basic authentication.

		Response (%)	Responses
Must have		16.67	2
Want		50.00	6
Don't care		33.33	4
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0





Additional Information (for the selected choice):

I have seen this implemented by providing a "test" URI which does not require authentication and then blocking access this URI using L7Policy.

Authentication for health monitoring is an edge case for us. It has never been asked for by our customers.

Sophisticated health check monitors that can support the gamut of HTTP features would be nice to have whether be it authentication, secure connection, POSTing a form, parsing responses, etc.

21) A project-user has a custom TCP-based service split across many application servers on the same subnet, and would like this service accessible via a single IP address. Further, he has written a custom load check for the back-end servers which powers custom auto-scaling logic. He would like to be able to adjust the weighting of any given back-end server in the load balancing pool on the fly, or add and remove some servers entirely. He needs to be able to do this without interrupting any TCP sessions still in progress.

		Response (%)	Responses
Must have		25.00	3
Want		16.67	2
Don't care		41.67	5
"DO NOT WANT"		16.67	2
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):



In my experience such logic is implemented using an orchestration system and not via LBaaS

This uses a mix of features, namely node member weighting and the ability to enable, drain and disable nodes.

This seems like a feature of whatever actually implements the load balancing rather than the API, but I can see the case for it and the API shouldn't exclude the possibility of it.

This would require standardizing how the load check is done with backend servers. Modern ADCs all have this in varying degrees, but there is no standard implementation, and therefore project-user would need to "know" which LB technology is going to be used, and would make his solution not portable.

22) A project-user has an HTTP-based load balanced service. His application has a wide range of completion times for each request, and would like the load balancers to always pick the back-end hosts with the fewest requests in progress for the next client request.




		Response (%)	Responses
Must have		58.33	7
Want		41.67	5
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

LEAST_CONNECTIONS algorithm. Absolute must have.

Bread and butter of loadbalancers is the variety of load-balancing algorithms, and "LEAST CONNECTIONS" algorithm is one of the basic ones, so yes a "must have" :)

23) A project-user has a very AJAX-intensive web application behind a load balancing service. He would like to speed up the performance for the web client by making sure the load balancer leaves the TCP connection between the web client and the load balancer open for a while after each request (so handshaking doesn't need to happen again). This application may be accessed via HTTP or HTTPS by the web client.

		Response (%)	Responses
Must have		25.00	3
Want		66.67	8
Don't care		8.33	1
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0





Additional Information (for the selected choice):

Is this policy per vip, per listener, per pool or per member?

We're talking keepalives here, I presume?

This is similar to HTTP Keep-Alive feature, and is a nice to have. Modern ADCs can enable this for an application, even when the backend application doesn't make use of Keep-Alives.

24) A project-user has an HTTP application behind a load balanced service. Some of the requests being processed take a very long time to complete, so she needs to make sure the load balancer doesn't close the connection with the client or server before processing is complete, and the content is returned from the back-end server servicing the request.





		Response (%)	Responses
Must have		33.33	4
Want		50.00	6
Don't care		8.33	1
"DO NOT WANT"		8.33	1
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

this is handled on many occasion by a global policy on the device. I would like to understand further the frequency of this

This is usually supported in ADCs by having a large client connection timeout value on the VIP.

25) A project-user has an HTTP application behind a load balanced service. In order to do A/B testing, the "A" back-end servers will be running entirely different code than the "B" servers, yet any web client connecting to the service needs to always be routed either to "A" or "B" hosts depending on which group they first connected to. The project-user would like both the "A" and "B" versions of the site accessible via the same IP address.

		Response (%)	Responses
Must have		8.33	1
Want		41.67	5
Don't care		41.67	5
"DO NOT WANT"		8.33	1
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):





this could be achieved via application logic and cookie based persistence

Just use a different IP address for the B test site.

Session persistence would accomplish this. However, our customers don't necessarily do this kind of testing (at least from the data I have).

This can be enabled by using a persistence mode on the VIP, so a client who gets load-balanced to an "A" server stays with an "A" server for the rest of the session.

26) A project-user has an HTTP application behind a load balanced service. In order to do maintenance on each back-end server, he would like to be able to sequentially remove servers from the load balancing pool, perform the maintenance, and add them back into the pool. He does not want to have to put the site into "maintenance mode" to do this, nor does he want any in progress requests getting interrupted because of this maintenance.

		Response (%)	Responses
Must have		41.67	5
Want		41.67	5
Don't care		8.33	1
"DO NOT WANT"		8.33	1
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

is this something like "don't remove member till all connection are done"?

Can't the admin just service apache stop on the node, do the maintenance, then service apache start? What does the LBaaS service have to do with this use case? This is the realm of configuration management tooling, not load balancers. The health check system of the LB will take any node out of rotation that fails the HTTP/80 ping check, so this is a non-use-case IMO



Enabling, draining and disabling nodes should allow for this use case.

Nice to have. Usually this is achieved by marking a backend-server state as a "draining" mode, meaning that the loadbalancer will stop sending traffic

to this backend server (apart from existing persistent connections), and wait for the server to "bleed off" all its existing persistent connections or finish processing existing requests. Once this is done, the backend server can be taken offline and patched up, then put back and its state changed to ACTIVE. Modern ADCs do support this.

or any TCP application.

27) A project-user has an HTTP application behind a load balanced service. In order to perform some major maintenance, he would like to be able to put the site into "maintenance mode" for the duration of the maintenance, and then take the site out of maintenance mode when the maintenance is over. In order to test the site before it goes live again, he would like to have a way to enable the site for specific client IPs only while in maintenance mode.

		Response (%)	Responses
Must have		0.00	0
Want		75.00	9
Don't care		25.00	3
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0




Additional Information (for the selected choice):

could be done same as use case 9 - I3 based policy

I'm between Don't care and Want on this one... basically, I can see the use in a total site maintenance mode being driven from the load balancer, but it's not something that is high priority to me.

Nice to be able to restrict access to a VIP to only a subset of client IP addresses (a "white list").

28) A project-user has an HTTPS site behind a load balanced service that powers many different e-commerce sites. Each of these e-commerce sites has its own SSL certificate and she would like the load balanced service to use the appropriate one when clients connect, according to the SNI protocol standard.




		Response (%)	Responses
Must have		33.33	4
Want		33.33	4
Don't care		33.33	4
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

Use a separate VM for load balancer for each site.

This is a popular feature in hosted environments, and as these move to cloud, this should be supported.

29) A project-user has an HTTP site behind a load balanced service that powers many different e-commerce sites. She has different groups of back-end servers that should be used for different sites, depending on the HTTP/1.1 hostname that the web client requests when connecting.

		Response (%)	Responses
Must have		25.00	3
Want		58.33	7
Don't care		16.67	2
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0



Additional Information (for the selected choice):

Use a separate VM for load balancer for each site.

Same as in use case #28, several websites on the same IP address (Apache virtual hosts) is very common.




30) A project-user has an HTTP application behind a load balanced service. This application has been attacked by malicious 3rd parties in the past coming from a specific block of IP addresses. She would like the load balancers to drop any requests from this block of IP addresses before they reach her back-end servers.

Response (%) Responses

Must have		50.00	5
Want		50.00	5
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0




Additional Information (for the selected choice):	
L3 source IP policies I think that such capabilities are delivered by another system and should also included detection and analysis capabilities	
I call this access list control.	
This is the opposite of use case #27 (a "black list")	

31) A project user wants to troubleshoot his application and needs to download a connection log from the load balance.

		Response (%)	Responses
Must have		54.55	6
Want		36.36	4
Don't care		9.09	1
"DO NOT WANT"		0.00	0
		Answered Question	11
		Skipped Question	1

Additional Information (for the selected choice):	
sometime it is preferable to push to a central syslog and allow the tenant to download from there	
Only Want instead of Must Have since a mostly similar (except without rejections on the LB side) thing can be gotten by logging centrally the application state logs.	
Absolute must have.	
Nice feature to be able to offer users.	

32) A project user wants to host a HTTP website which includes a Java applet. The applet requires an additional TCP port to communicate with its backend application. The Java security policy for applets restricts them to only connecting to the host they were retrieved from.

		Response (%)	Responses
Must have		18.18	2
Want		27.27	3
Don't care		54.55	6
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):	
support for an additional listener on non default tcp ports	
is this a use case? I don't see a description of what the end user would be doing or needing? Does anyone even use Java applets anymore?	
Interesting use case for multiple ports.	
Java applets are a dying client-side technology, so this is not an important use case for me.	

33) A project-user has a 2048-bit wildcard SSL certificate for "*.example.com" and a 4096-bit SSL certificate for "secure.example.com." She wants to serve all her websites that apply to the "example.com" domain from the same IP (using the SNI standard), and also wants to make sure that any requests for "secure.example.com" use the 4096-bit certificate.



		Response (%)	Responses
Must have		8.33	1
Want		75.00	9
Don't care		16.67	2
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

could this also be handled by ordering of certificates? I suspect this use case is not very frequent.

An advanced use case of flexible SSL certificate handling. Don't see much use of this in practice.

34) A project-user has an old 2048-bit wildcard SSL certificate for "*.example.com" and a new 4096-bit wildcard SSL certificate for "*.example.com" from a new certificate authority. She wants all of the various hostnames that apply to the "example.com" domain served from the same IP (using the SNI standard), and usually using the old certificate-- but she wants requests for "secure.example.com" and "admin.example.com" to use the new certificate.



		Response (%)	Responses
Must have		8.33	1
Want		83.33	10
Don't care		8.33	1
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

could this also be handled by ordering of certificates?

same comment as use case #33

35) A project-user has a web application that contain unsecured and secured parts (ex: unsecured - http://www.acme.com/app/catalog and secured https://www.acme.com/checkout). When an application user access the URI that is handled by the secured part via unsecured protocol (ex: http://www.acme.com/checkout), she will be redirected to the secured section (ex: https://www.acme.com/checkout) and when an application user access the URI that is handled by the unsecured part via a secured protocol (ex: https://www.acme.com/app/catalog) she will be redirected to the unsecured part (ex: http://www.acme.com/app/catalog)

		Response (%)	Responses
Must have		0.00	0
Want		58.33	7
Don't care		41.67	5
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0




Additional Information (for the selected choice):

this is an example where L7 policy reuse between listeners makes configuration simpler while the action should be out of the policy

This seem like it could/should be implemented in the application using the X-Proto header.

Nice to have. Part of basic L7 content processing features.

36) A project-user consistently wishes to classify static content so that she can use the same classifications for all of it's web applications. For this she wishes to classify all picture types as "pictures", classify all streaming types as "videos" and classify all word office documentation type as "word-documents". She wants to reuse this definition to direct traffic to the appropriate static content servers. For example browsing to http://www.acme.com/uri1/piv.gif will be handled by the static content server hosting images appropriate for www.acme.com and browsing to http://www.mysite.com/myname/portrait.jpg will be handled by the static content server hosting images appropriate for www.mysite.com

		Response (%)	Responses
Must have		16.67	2
Want		50.00	6
Don't care		33.33	4
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

this is an example of system wide l7 policy reuse

Nice to be able to reuse of existing definitions, but I don't see this as a must have,

37) A project user wants to terminate SSL connections so that the load balancing logic will be able to make decision based on layer 7 information. He then wishes the traffic to continue to be encrypted between the load balancer and the application servers. He does not care to provide trusted certificates

		Response (%)	Responses
Must have		58.33	7
Want		25.00	3
Don't care		16.67	2
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

this is same as use case 8

Our customers have been asking for the re-encryption feature for a while, but mostly for X-Forwarded-For header.

This is an important feature in clouds to support a range of applications in regulatory environments such as healthcare or finance.

38) A project user wants to specify the allowed cipher suites and the allowed SSL protocols when terminating a connection on the load balancer. He only want to allow TLS1.2 and Elliptic curve cyphers to be used.

		Response (%)	Responses
Must have		16.67	2
Want		50.00	6
Don't care		25.00	3
"DO NOT WANT"		8.33	1
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

I would like to be able to specify the TLS ciphers supported, but not necessarily Elliptic curve ciphers.

Don't think that end users should be able to dictate what protocols to use. This should be part of the operator setup and deployment requirements.

39) A project user wants to manage the allowed protocols and cipher suites in a central way so that if it needs modification dues to new security concerns, it will affect all applications.

		Response (%)	Responses
Must have		8.33	1
Want		41.67	5
Don't care		41.67	5
"DO NOT WANT"		8.33	1
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

an example of SSL policy reuse

Same comment as #39

40) A project user has an HTTPS service on the load balancer that has has a back-end pool of servers which should also be accessed over HTTPS. The user is not very sophisticated (ie. doesn't want to have to manage a whole PKI). Each of these back-end servers is configured with its own separate self-signed server certificate, and the user would like the load balancer to authenticate that the self-signed cert is the one configured for that server in order to protect against potential man-in-the-middle attacks.



		Response (%)	Responses
Must have		33.33	4
Want		33.33	4
Don't care		33.33	4
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

support for trusted certificates for re encryption

See this as an advanced capability of SSL offload. Not much use in the real world.

41) A project user has several load balanced services and would like to collect rolling data on bytes transferred in and out, number of connections and other aggregate statistics per service. Ideally, this would be data that could be collected via sensu or some other means suitable for generating a dashboard graph.



		Response (%)	Responses
Must have		41.67	5
Want		58.33	7
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

discuss if this is an Ibaas, ceilometer or both

Nice feature to offer users, so they can learn about traffic patterns for their app.

42) A project user has a load balanced service and would like to see bytes transferred, connections, and other basic statistics on each pool of servers in the back-end.




		Response (%)	Responses
Must have		25.00	3
Want		75.00	9
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):



a more in-depth analysis should be done on the required for monitoring metrics

Yes this is the least "basic" of statistics that we can offer, and this is a must.

43) L7 content switching

		Response (%)	Responses
Must have		41.67	5
Want		50.00	6
Don't care		8.33	1
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

44) Session persistence

		Response (%)	Responses
Must have		75.00	9
Want		25.00	3
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

45) Health Monitoring

		Response (%)	Responses
Must have		90.91	10

Want		9.09	1
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

What is this? There needs to be more details on this one.

46) IPv4 and IPv6 support

		Response (%)	Responses
Must have		66.67	8
Want		33.33	4
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

47) SSL termination

		Response (%)	Responses
Must have		91.67	11
Want		8.33	1
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

48) SSL back end encryption

		Response (%)	Responses
Must have		45.45	5
Want		54.55	6
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

Needs more detail... this seems to have been covered by some use cases above.

49) SNI support

		Response (%)	Responses
Must have		50.00	5
Want		50.00	5
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	11
		Skipped Question	1

Additional Information (for the selected choice):

Needs more detail... this seems to have been covered by some use cases above.

50) UDP support

		Response (%)	Responses
Must have		9.09	1

Want		63.64	7
Don't care		27.27	3
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):
 Needs more detail... this seems to have been covered by some use cases above.

51) Round Robin algorithm support

		Response (%)	Responses
Must have		100.00	11
Want		0.00	0
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):
 Needs more detail... this seems to have been covered by some use cases above.

52) Backup servers

		Response (%)	Responses
Must have		54.55	6
Want		45.45	5
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):
 Needs more detail... this seems to have been covered by some use cases above.

53) Spillover

		Response (%)	Responses
Must have		0.00	0
Want		50.00	5
Don't care		50.00	5
"DO NOT WANT"		0.00	0
		Answered Question	11
		Skipped Question	1

Additional Information (for the selected choice):
 Needs more detail... this seems to have been covered by some use cases above.
 I'm not sure what "Spillover" means.

54) Max Pool/Member Connections

		Response (%)	Responses
Must have		45.45	5
Want		54.55	6
Don't care		0.00	0
"DO NOT WANT"		0.00	0
		Answered Question	12
		Skipped Question	0

Additional Information (for the selected choice):

Needs more detail... this seems to have been covered by some use cases above.

55) Domain based pool members

	Response (%)	Responses
Must have	0.00	0
Want	81.82	9
Don't care	18.18	2
"DO NOT WANT"	0.00	0
	Answered Question	12
	Skipped Question	0

Additional Information (for the selected choice):

is this using L7 policy to select group based on hostname / domain name?

Needs more detail... this seems to have been covered by some use cases above.

I am worried that domain-based pool members can significantly slow down load balancing operations. However, as I don't anticipate many of our customers will use this feature, it's not a show-stopper for us to have it in the product.