VIRTUALIZATION HEALTH CHECK

Abstract

A health check using Foglight Evolve has been undertaken to: reveal resource tuning requirements that both improve performance and regain resources for future use, and; check on current and forecast capacity to aid planning and avoid contention, and; check for serious or frequently occurring problem that might impact on applications.

Contents

Management Summary

Foglight Evolve has been used to review the health of the virtual infrastructure at ABC Company. The review has concentrated on three main areas:

- Optimization opportunities that both improve performance and facilitate the recovery of resources for future use;
- Current and forecast capacity to aid planning and avoidance of contention;
- Frequently occurring and / or serious problems that might impact on applications

The application was deployed at the end of September to collect and analyse data. The reports and screenshots cover the period between the 2nd and the 12th of October.

Optimization: Foglight has uncovered potential for up to **$45,444** of savings in the form of deferred costs of further hardware acquisition.

Capacity: Foglight has determined that at a high level there is sufficient capacity available to support a further 43 virtual machines similar to the current average resource consumptions. The first constraint that keeps this number from being larger is memory. At current trend this capacity is sufficient for 232 days.

Monitoring: Though a review of individual virtual machine performance using CASK Theory shows only modest negative impact, the alarm review show lots of latency events particularly at the storage level, many of which are of lengthy duration.

Alarm Source	Alarm Count	Severities			Alarm Duration		
		F	C	W	Min	Max	Avg
VMW Datastore Average Latency	1.156	468	305	383	-51,2 min	7,8 hr	20 min
StSAN E Physical Disk Busy w Email	1.148		378	770	15 min	2,7 hr	1,1 hr
VMW ESX Server Total Latency	600	167	93	340	-51,2 min	7,8 hr	23 min

Optimization

Foglight has uncovered potential for up to **$45,444** of savings in the form of deferred costs of further hardware acquisition. The breakdown can is shown below:

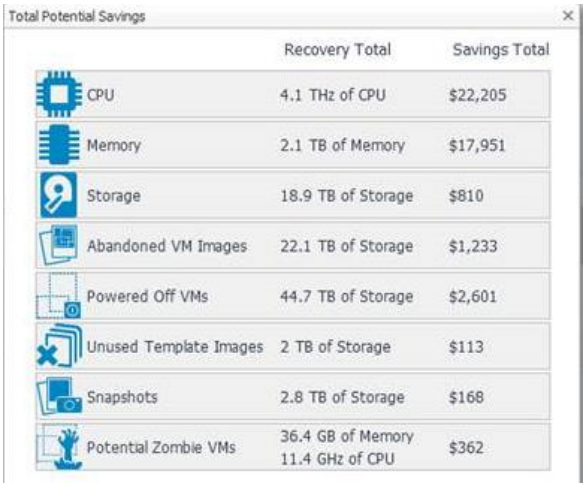

Total Potential Savings $45,444

Calculations within the optimizer analytics are governed by setting which determine how aggressively (or otherwise) recommendations for optimization and right-sizing should be applied. The defaults are intended to provide for the maximum utilizations yet seen plus and overhead for growth. The settings used were as shown below:

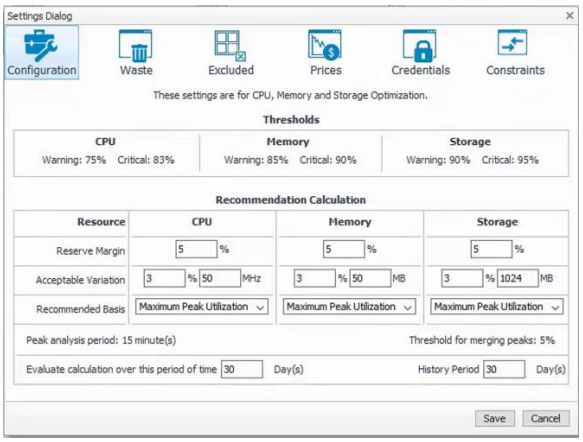

CPU Optimization

Foglight Evolve's optimization analytics has determined that there are potential deferred cost **savings for CPU of $22,205**. There are more than 120 individual recommendations for decreasing CPU allocations, increasing allocations for those virtual machines at or close to their current limits, changing reservations and removing limits. A variety of these are shown below:

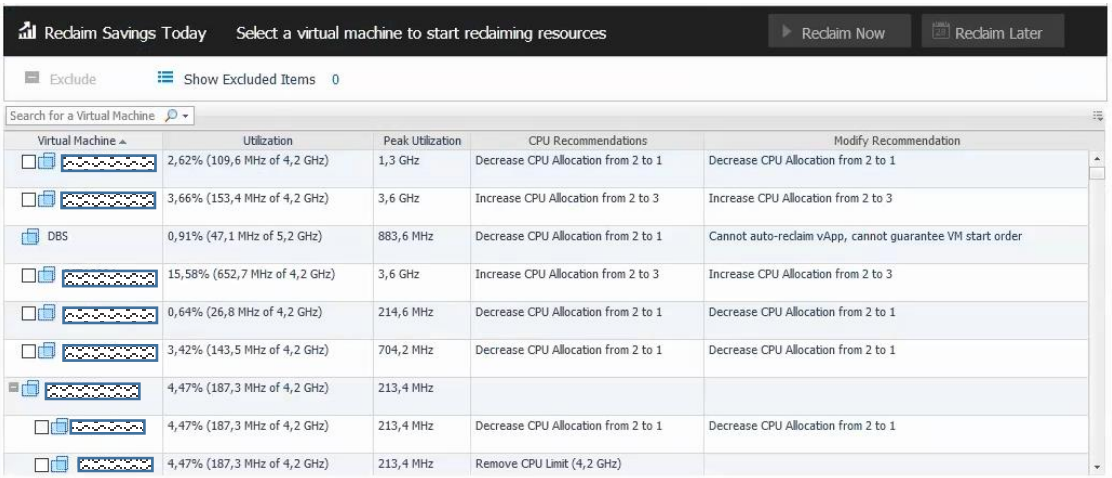

Virtual Machine	Utilization	Peak Utilization	CPU Recommendations	Modify Recommendation
[redacted]	2,62% (109,6 MHz of 4,2 GHz)	1,3 GHz	Decrease CPU Allocation from 2 to 1	Decrease CPU Allocation from 2 to 1
[redacted]	3,66% (153,4 MHz of 4,2 GHz)	3,6 GHz	Increase CPU Allocation from 2 to 3	Increase CPU Allocation from 2 to 3
DBS	0,91% (47,1 MHz of 5,2 GHz)	883,6 MHz	Decrease CPU Allocation from 2 to 1	Cannot auto-reclaim vApp, cannot guarantee VM start order
[redacted]	15,58% (652,7 MHz of 4,2 GHz)	3,6 GHz	Increase CPU Allocation from 2 to 3	Increase CPU Allocation from 2 to 3
[redacted]	0,64% (26,8 MHz of 4,2 GHz)	214,6 MHz	Decrease CPU Allocation from 2 to 1	Decrease CPU Allocation from 2 to 1
[redacted]	3,42% (143,5 MHz of 4,2 GHz)	704,2 MHz	Decrease CPU Allocation from 2 to 1	Decrease CPU Allocation from 2 to 1
[redacted]	4,47% (187,3 MHz of 4,2 GHz)	213,4 MHz		
[redacted]	4,47% (187,3 MHz of 4,2 GHz)	213,4 MHz	Decrease CPU Allocation from 2 to 1	Decrease CPU Allocation from 2 to 1
[redacted]	4,47% (187,3 MHz of 4,2 GHz)	213,4 MHz	Remove CPU Limit (4,2 GHz)	

Memory Optimization

Foglight Evolve's optimization analytics has determined that there are more than 30 rightsizing activities required but that the net result of these regarding deferred costs / savings is neutral. Whilst cost neutral they will improve overall performance. The 33 individual recommendations include decreasing memory allocations, increasing allocations for those virtual machines at or close to their current limits and removing some imposed limits. A variety of these are shown below:

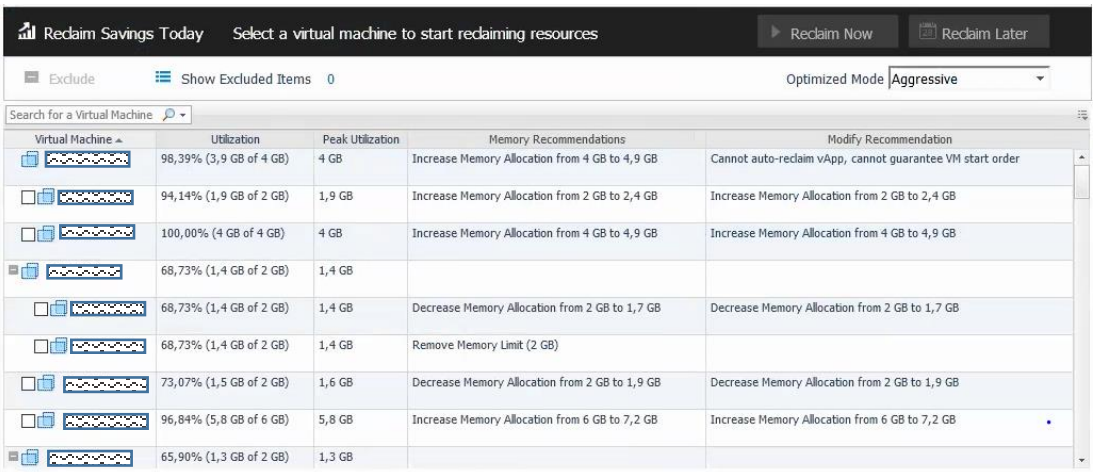

Virtual Machine	Utilization	Peak Utilization	Memory Recommendations	Modify Recommendation
[redacted]	98,39% (3,9 GB of 4 GB)	4 GB	Increase Memory Allocation from 4 GB to 4,9 GB	Cannot auto-reclaim vApp, cannot guarantee VM start order
[redacted]	94,14% (1,9 GB of 2 GB)	1,9 GB	Increase Memory Allocation from 2 GB to 2,4 GB	Increase Memory Allocation from 2 GB to 2,4 GB
[redacted]	100,00% (4 GB of 4 GB)	4 GB	Increase Memory Allocation from 4 GB to 4,9 GB	Increase Memory Allocation from 4 GB to 4,9 GB
[redacted]	68,73% (1,4 GB of 2 GB)	1,4 GB		
[redacted]	68,73% (1,4 GB of 2 GB)	1,4 GB	Decrease Memory Allocation from 2 GB to 1,7 GB	Decrease Memory Allocation from 2 GB to 1,7 GB
[redacted]	68,73% (1,4 GB of 2 GB)	1,4 GB	Remove Memory Limit (2 GB)	
[redacted]	73,07% (1,5 GB of 2 GB)	1,6 GB	Decrease Memory Allocation from 2 GB to 1,9 GB	Decrease Memory Allocation from 2 GB to 1,9 GB
[redacted]	96,84% (5,8 GB of 6 GB)	5,8 GB	Increase Memory Allocation from 6 GB to 7,2 GB	Increase Memory Allocation from 6 GB to 7,2 GB
[redacted]	65,90% (1,3 GB of 2 GB)	1,3 GB		

Storage Optimization

Foglight Evolve's optimization analytics has determined that there are potential deferred cost **savings for Storage of $810**. There are more than 150 individual recommendations relating to the ABC Company-Group cluster alone. These recommendation mostly relate to the recovery of over allocated storage and for the ABC Company-Group cluster this would result in reclamation of 2.8TB. Examples of these recommendations can be seen below:

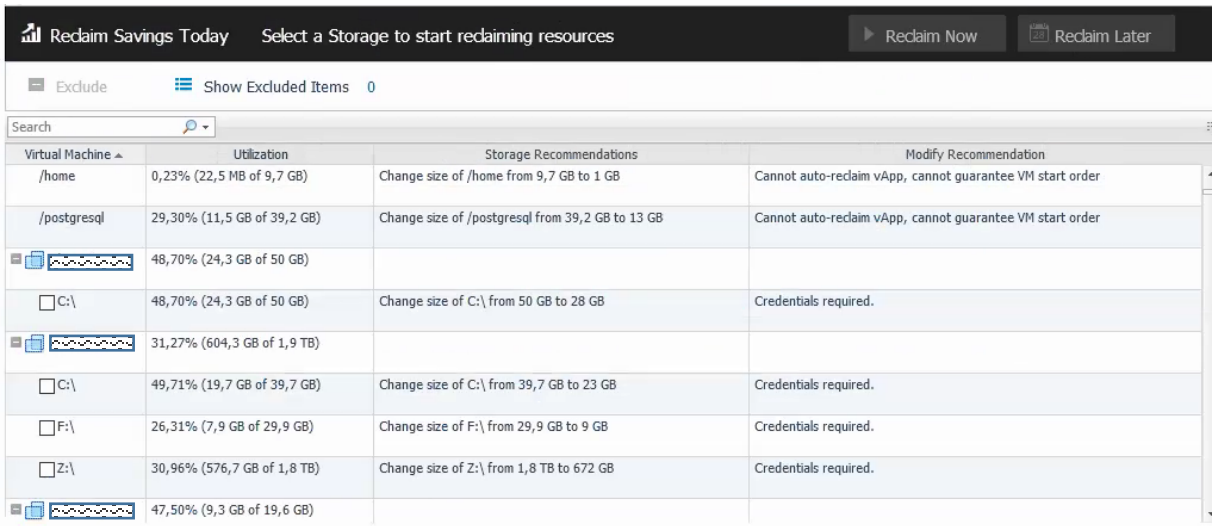

Virtual Machine	Utilization	Storage Recommendations	Modify Recommendation
/home	0,23% (22,5 MB of 9,7 GB)	Change size of /home from 9,7 GB to 1 GB	Cannot auto-reclaim vApp, cannot guarantee VM start order
/postgresql	29,30% (11,5 GB of 39,2 GB)	Change size of /postgresql from 39,2 GB to 13 GB	Cannot auto-reclaim vApp, cannot guarantee VM start order
[redacted]	48,70% (24,3 GB of 50 GB)		
C:\	48,70% (24,3 GB of 50 GB)	Change size of C:\ from 50 GB to 28 GB	Credentials required.
[redacted]	31,27% (604,3 GB of 1,9 TB)		
C:\	49,71% (19,7 GB of 39,7 GB)	Change size of C:\ from 39,7 GB to 23 GB	Credentials required.
F:\	26,31% (7,9 GB of 29,9 GB)	Change size of F:\ from 29,9 GB to 9 GB	Credentials required.
Z:\	30,96% (576,7 GB of 1,8 TB)	Change size of Z:\ from 1,8 TB to 672 GB	Credentials required.
[redacted]	47,50% (9,3 GB of 19,6 GB)		

Abandoned Virtual Machine Images

Foglight identified no abandoned virtual machine images in the monitored environment.

Powered Off Virtual Machines

Foglight identified a single virtual machine in the monitored environment that had been powered off for more than 30 days.

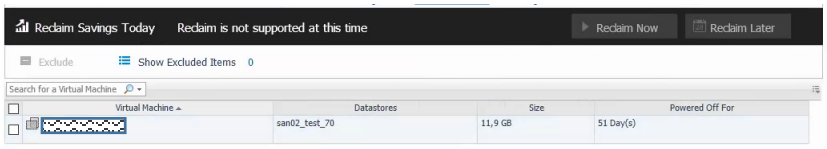

Unused Template Images

Foglight identified no unused template images in the monitored environment.

Orphaned and Outdated Snapshots

Foglight identified no orphaned or out of date snapshots in the monitored environment.

Potential Zombie Virtual Machines

Foglight Evolve identified 20 virtual machines whose usage pattern suggests that they are running, consuming resources but not actively being used. It is possible that this categorization might be incorrect but each should be investigated both to reclaim the resources and to improve the performance of the remainder of the environment. Examples are shown below, note that 2 have already been investigated and excluded from the list of potential zombies.

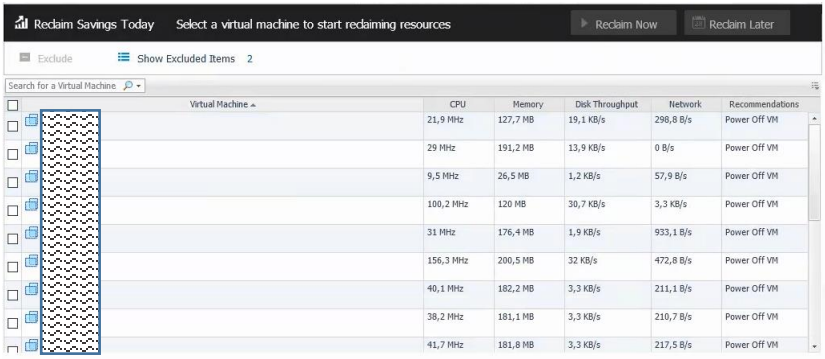

Realizing the Value of Recommendations

It should be noted that a number of tools provide recommendations but that rarely do organizations manage to realise the value they promise. The reason for this can be illustrated by a worked example regarding practicality.

Imagine that there were 100 recommendations awaiting action. Each action requires opening the vSphere UI, finding the appropriate object (VM), navigating to the correct point and then making the change. Let's assume that for each recommendation this process can be accomplished in 6 minutes (this is typically an under estimate).

100 recommendations x 6 minutes = 600 minutes = 10 hours

That is 10 hours of solid work with no breaks, no interruptions, etc. This means that the work would need to be spread over several days. Moreover some of the operations can only be completed when the underlying systems can be shutdown for maintenance. All this means that in practice only a small number of recommendations are acted upon before people give up.

Foglight Evolve is different. Foglight has an embedded workflow automation system. To take action on a recommendation, the user selects those they wish to apply (without changing to another user interface) and then chooses to take action "Now" (for those

things that do not require a reboot) or to schedule them to occur “Later” during a maintenance window.

The user can go further once they trust the recommendations in a particular category whereupon they can configure the system to take care of them entirely automatically during maintenance periods without presenting them or requiring manual intervention.

Foglight Evolve makes it possible to realize the value that others promise and fail to deliver.

Capacity

Foglight Evolve has determined that at a high level there is sufficient capacity available to support a further 43 virtual machines similar to the current average resource consumptions. The first constraint that keeps this number from being larger is memory. At current trend this capacity is sufficient for 232 days.

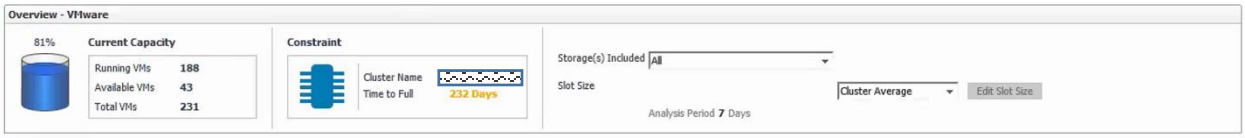

This is further broken down by cluster. We can see that the primary cluster is the ABC Company-Group cluster consisting of 4 ESX Hosts and currently running 167 virtual machines. Foglight has calculated the mean size of the existing VMs in each resource category and has determined that there is capacity on this cluster for a further 40 VMs matching this average. The first resource constraint preventing a greater number being deployed is Memory.

Foglight has analysed the resource consumption and predicted forward usage.

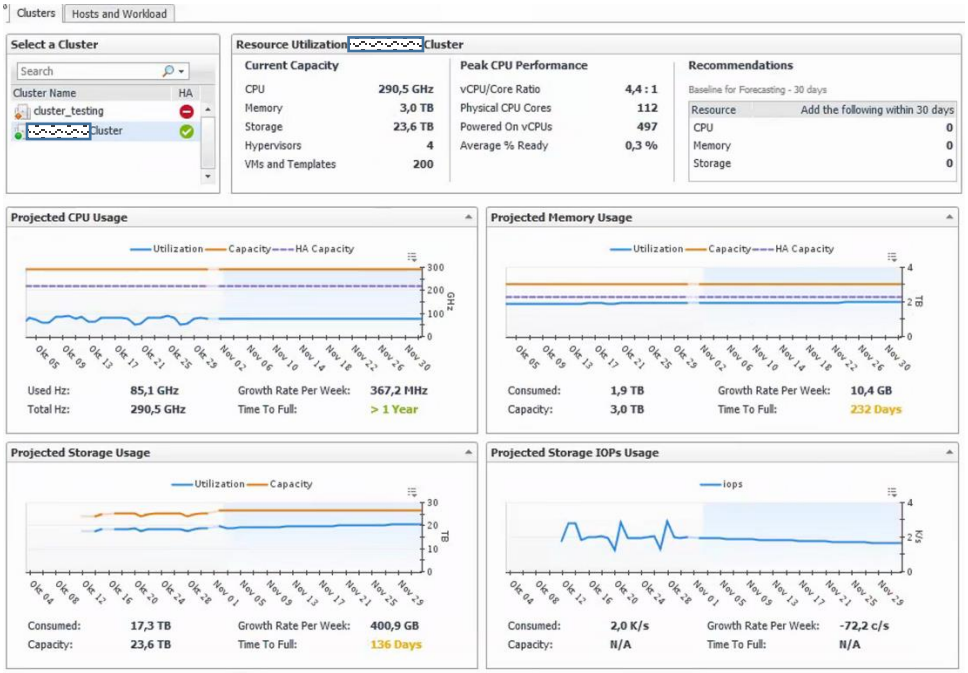

Monitoring

Foglight Evolve has analysed individual virtual machine performance using CASK Theory. This essentially says that a virtual machine is most constrained by its most constrained resource.

CASK analysis shows a very bad score for [redacted] which is suffering from severe disk latency (higher scores being good and lower ones bad).

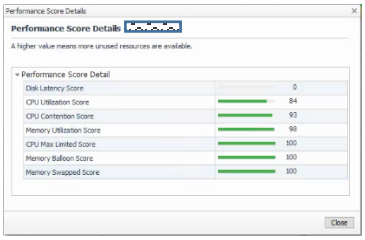

This situation is corroborated by the Alarm Analysis which shows lots of latency events particularly at the storage level both at the data store and physical storage layers, many of which are of lengthy duration.

Alarms by Source

Search

Alarm Source	Alarm Count	Severities			Alarm Duration		
		F	C	W	Min	Max	Avg
VMW Datastore Average Latency	1.156	468	305	383	-51,2 min	7,8 hr	20 min
StSAN E Physical Disk Busy w Email	1.148		378	770	15 min	2,7 hr	1,1 hr
VMW ESX Server Total Latency	600	167	93	340	-51,2 min	7,8 hr	23 min
Catalyst Garbage Collector Check	386			386	7 sec	39 sec	9 sec
StSAN E LUN Total Latency over Threshold w Email	341		279	62	15 min	21 d	1,5 d
StSAN E Filer Volume Write Latency Over Threshold w Email	280		258	22	15 min	1,3 hr	18 min
VMW PNICs Packet Loss	134			134	5,0 min	1,6 hr	8,4 min
VMW Standard Virtual Switch Network Packet Loss	133			133	5,0 min	1,6 hr	8,3 min
StSAN E Filer Volume Read Latency over Threshold w Email	119		72	47	15 min	1,7 hr	23 min
VMW Virtual Machine Balloon Memory Deflation	73		59	14	5,0 min	18 d	9,9 hr
VMW Virtual Machine Memory Swapping	71	5	66		-51,3 min	5,2 hr	27 min
VMW ESX Server Queue Latency	62	14	11	37	4,8 min	50 min	12 min
VMware Virtual Machine OS reboot	58	58			-1,8 hr	5,3 hr	18 min
VMW Datastore Inactive	56	56			5,0 min	18 d	1,7 d
VMW Virtual Machine VMware Tools	55			55	4,1 min	21 d	11 d
VMW Virtual Machine Logical Drive Estimated Fill Time	48		16	32	24 hr	6,0 d	1,7 d
VMW ESX Server Balloon Memory Deflation	46		28	18	-51,3 min	7,8 hr	1,3 hr
VMW Virtual Machine Memory Utilization	44	9	8	27	4,8 min	20 min	6,6 min
VMW Virtual Machine Logical Drive Availability	38			38	-50,9 min	18 d	3,0 d
StSAN E NASVolume Low Available Capacity with Email	24		9	15	21 d	21 d	21 d
VMW ESX Server Memory Utilization Upward Trend	21	21			60 min	3,0 hr	1,7 hr

The two example alarms shown below are typical of the many:

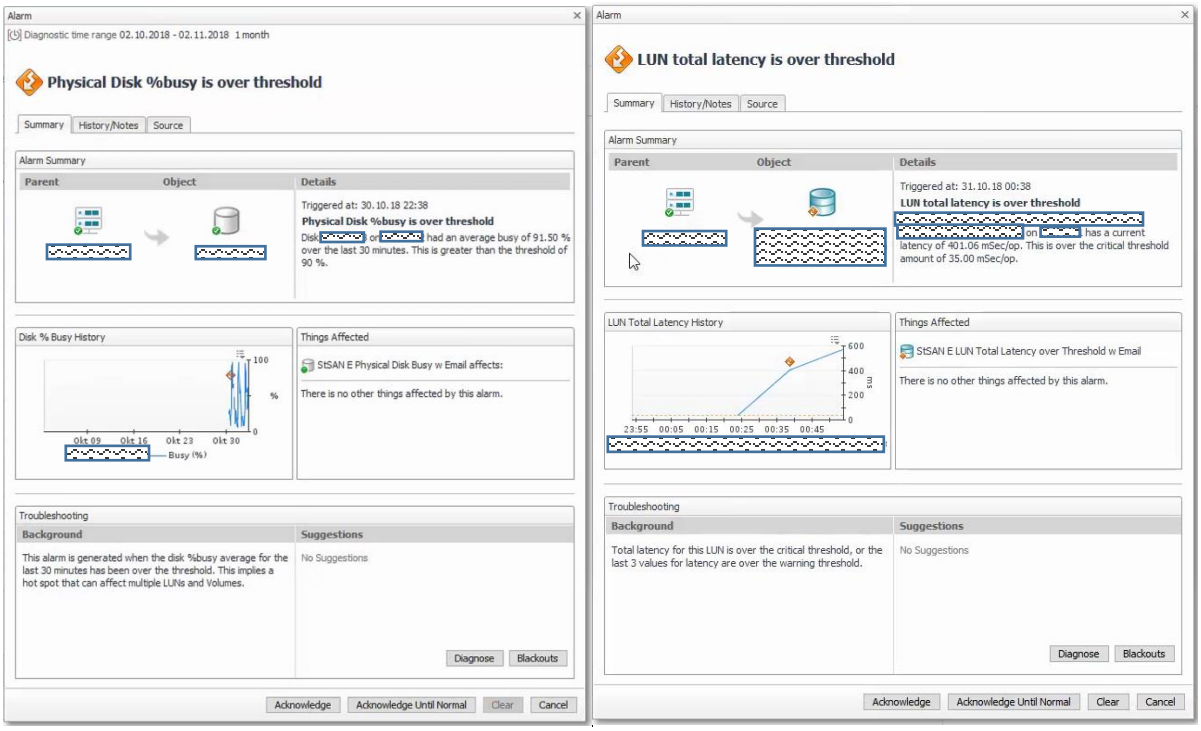

For more information, please contact your Quest Software Account Representative, or visit quest.com/Foglight-Evolve.