BROADBAND TESTING

Kemp Application Experience (AX) Report

A Broadband-Testing Report

First published May 2019 (V1.0)

Published by Broadband-Testing

E-mail : info@broadband-testing.co.uk
Internet : HTTP://www.broadband-testing.co.uk

TABLE OF CONTENTS

BROADBAND-TESTING

Broadband-Testing is an independent testing operation, based in Europe. Broadband-Testing interacts directly with the vendor, media, analyst, consultancy and investment communities equally and is therefore in a unique space within IT.

Testing covers all aspects of a product/service from business rationalisation in the first instance to every element – from speed of deployment and ease of use/management, through to performance and accuracy.

Testing itself takes many forms, from providing due diligence for potential investors through to public domain test reports.

Broadband-Testing is completely vendor neutral and independent. If a product does what it says on the tin, then we say so. If it doesn't, we don't tell the world it does what it cannot do. The testing is wholly complementary to analyst-related reports; think of it as analysts getting their hands dirty by actually testing what's on the latest hype curve to make sure it delivers on its claims and potential.

Broadband-Testing operates an **Approvals** scheme which prioritises products to be short-listed for purchase by end-users, based on their successful approval, and thereby short-cutting the evaluation process.

Output from the testing, including detailed research reports, articles and white papers on the latest IT technologies, are made available free of charge on our web site at HTTP://www.broadband-testing.co.uk

EXECUTIVE SUMMARY

- The load-balancer market and its requirements have changed significantly over the past few years; the original focus on hardware-driven solutions has been overtaken by hybrid hardware/virtual/cloud-based architectures.

- With the current trend of migrating parts of IT to the cloud, or multiple clouds, while maintaining some OnPrem element, a contemporary load-balancer solution has to be flexible, capable of supporting multiple, concurrent data and application sources, as well as multiple 3rd party platforms.

- Equally importantly, these environments must be manageable from a single platform and from essentially anywhere. The ability to see the "bigger picture" from top-down, then drill-down into specifics, is crucial to ensuring a truly optimised working architecture.

- Similarly, optimisation, analysis and security should not be three individual solutions fighting for IT resource and expertise; integration of the three is fundamental to successful load balancing, as is automation thereof as much as is possible.

- As the world moves from a data-centric view to an application-based approach – and not before time – being able to optimise and manage the "network" from an app perspective is increasingly critical. With its "AX" – Application Experience – approach to load balancing, Kemp appears to have fully grasped this requirement.

- Managing and optimising IT is no longer about bits and bytes and hardware throughput; data and applications can be located anywhere, on physical and/or virtual, or cloud-native platforms, but have to be treated as one giant resource. Equally, that resource must be able to effectively scale exponentially and infinitely. Gone are the days when it was (just about) possible to plan capacity and performance requirements on a two to five-year basis. At best, such an approach often led to expensive, over-engineered solutions in order to guarantee application and data availability, or led to costly downtime and unavailable resource. So, a traditional approach to optimising and managing load balancing is simply no longer an option.

- Load balancing is not just about optimisation but also about controlling costs and automating many administrative tasks helps enormously here. Also, many companies have already had their fingers badly burned by assuming that moving to the cloud means lower costs – it doesn't, unless properly managed. With the whole "food chain" from DevOps through to day to day application management now increasingly managed as a single IT entity, the ability to both scale on demand AND know the cost implications of this scaling is an essential marriage, in order to avoid budget nightmares. Again, with its metered-licensing, Kemp seems to have grasped this need; resource demands can rise or fall and costs will be representative of the demand, not fixed.

THE NEW WORLD OF LOAD BALANCING

The impact of the cloud and the at least partial move to OffPrem outsourcing, combined with a very different form of load balancing, has rewritten the IT rulebook in many ways.

Not least of these is in relation to load balancers themselves. Historically, the location of applications and data was relatively fixed; OnPrem on servers or residing in a monolithic outsourced or centralised DC (Data Centre). But now, that application and data source might be spread, not just between On/OffPrem, but also between multiple CSPs (Cloud Service Providers) who themselves have massively distributed DC architectures.

At the same time, what was once a very physical environment – real servers, traditional storage – is now very much virtualised, with applications running in VMs (Virtual Machines) or containers – or both, with the load balancing element itself also often sitting as a virtual application. This has resulted in a radically different architecture and data/application hierarchy to manage and optimise delivery of.

It doesn't take an IT architecture specialist to therefore work out that solutions designed in the early 2000s to manage load balancing back then are not fit for purpose as we approach the third decade of the century. With most enterprise users seeing SaaS (Software as a Service) based applications as their now established delivery format, whether developed in-house or from software vendors, even those who haven't yet moved to the cloud have a significant OffPrem presence to their IT landscape.

From a load balancing vendor perspective, the initial response of virtualising their (previously hardware-only) appliances was a step in the right direction, certainly, but it doesn't solve the problem of managing the new IT environment, where it's not simply a case of load balancing globally, but effectively delivering anything from anywhere, in the same manner, and from a single point of management. Trivial? Certainly not.

Kemp – the subject of this report – describes the new landscape as requiring an "always-on" application experience (AX), marketese for sure, but we get what they are defining – continuously optimising application availability - and this is applicable all the way from DevOps to end users. And it's not just about optimising delivery, but securing it too. One major issue for IT administrators in recent years has been how to whittle down the plethora of optimisation and security products and technologies available to a manageable and affordable combination. So, the more capability one platform has to optimise, orchestrate and secure application delivery, the fewer the headaches, not to mention the budgetary requirements.

Load Balancing - Before and After…

Historically, in our "before" scenario, a load balancing solution was delivered as a physical or (more recently) virtual appliance, or indeed as a pair for HA (High Availability), with the load balancer sat in front of the application to optimise its access and availability. This included providing security features in the form of SSL offload and a web application firewall to protect applications from cyber-attacks.

One primary issue with this traditional approach is that, in rolling out appliances, a certain amount of guesswork is required in terms of required capacity over the next, say, two to five years. The natural conclusion of this form of "guesstimation" is either expensive over-

engineered solutions, in order to absolutely guarantee availability at all times, or – alternatively - damaging under-provisioning, where you run out of capacity, probably at the most critical timepoints. Under-budgeting can also lead to a potentially catastrophic single point of failure and a resulting wide area of fallout, in terms of applications and user base being impacted upon, and that means your business suffers.

At a more fundamental level, this approach simply isn't in line with the contemporary DevOps toolchain or a continuous implementation-continuous delivery (CI/CD) model. Unsurprising, therefore, that ZK Research found that, out of 1,000 enterprise survey respondents, 70% wanted to replace their current load balancing vendor, citing vendor lock-in, pricing and unnecessary complexity as key reasons. Kemp itself has been analysing the key problem areas of load balancing operations, and found that, from in excess of 60,000 Kemp global application deployments, 61% of application availability related events were caused by application capacity issues such as failing or degrading application servers.

The "after" scenario sees not only the guesswork replaced by capacity on demand and total management of an entire estate, regardless of where those applications live, but also – with automation – the ability to be proactive in the above case, so server issues would no longer have the same level of impact upon application availability. According to research from Enterprise Management Associates (EMA), 92 percent of enterprise network engineering and operations professionals have an initiative in place to expand their use of network automation, and 70 percent say this automation is a high priority. In light of the aforementioned capacity issues, it is easy to see why.

Let us now look in more detail at what Kemp is offering, followed by some hands-on examples of the technology in action.

PRODUCT OVERVIEW

Kemp's Application Experience (AX) has been developed with the aim of optimising performance, security and application availability on a 24x7 basis.

The AX fabric provides all components from licensing, through deployment, to management and monitoring. The AX solution is based around three core components: The LoadMaster appliances (physical and virtual/cloud), the Kemp 360 Central management/orchestration interface and the Kemp 360 Vision analytics.

Kemp LoadMaster Appliances

The appliance range is basically formed of two sub-ranges, all of which interoperate – the hardware and the virtual-based instances, the latter also supporting a number of CSPs.

Hardware options include an entry-level LoadMaster X3 with 3.4Gbps throughput and 1,700 SSL tps, through the X15 with up to 15Gbps of Layer 7 throughput, up to the LoadMaster X40 which provides 40Gbps throughput and up to 35,000 SSL tps and over 75 million concurrent Layer 4 connections. All the hardware devices can be clustered into a single, manageable entity, while providing redundancy at the same time.

Figure 1 – Kemp LoadMaster X-15

The Edge Security Pack (ESP) enables pre-authentication of clients (LDAP, Active Directory, RADIUS and SAML) while the Web Application Firewall (WAF) offers continuous application level protection through daily rule updates.

Kemp's Virtual LoadMaster (VLM) is still fully featured, including SSL offload. It can be deployed on all the major hypervisor platforms and on the leading public cloud services. Key to the deployment is that the feature set is consistent, regardless of where it is deployed. At the entry level, the VLM range offers 200Mbps throughput and 200SSL tps, moving up to 10Gbps throughput and 12,000 SSL tps at the top of the range.

Kemp 360 Central

Kemp has a single management console for the entire deployment, in the form of Kemp 360 Central, which provides orchestration and control for LoadMaster and 3rd party load balancers across cloud, hypervisor and hardware platforms. The idea is that you can manage the estate from anywhere, including support for 3rd party vendor appliances such as F5/NGINX, HAproxy and AWS ELB. It also extends to support over multiple platforms, including public cloud and local hypervisors. Many of the tasks can be automated, thereby reducing time and associated costs. All event logging can be consolidated, and a single dashboard shows the real-time status of the estate, regardless of location and appliance type. The aim is to provide instant visibility of the health of all application delivery resources from a single point. From a deployment perspective, settings can be pre-configured and deployed onto any of the supported platforms from Central, including virtual and cloud-native load balancing instances across cloud and hypervisor platforms. Using templates is a guarantee of error-free provisioning load balancing instances and application services. Given that around 80% of all network problems are actually human configuration errors, this is significant. By collating events and metrics, Central is also designed to rapidly identify issues across the load balancing infrastructure.

Kemp 360 Vision

The analytics of the AX solution is expanded through Kemp 360 Vision – a GUI-based application that correlates observed events and conditions across all load balancing resources, with the aim of identifying and remediating issues. It does this through the use of predictive analytics that provide early warning of emerging conditions. With Vision you can centrally monitor application performance and availability across all public and private cloud infrastructures, with in-context issue analysis and actionable alerts delivered through communications channels such as Slack, SMS or email.

In addition to the Kemp real estates, Vision monitors the complete application environment, including 3rd party appliances, so it is largely a one-stop shop solution, minimising the need for multiple interfaces and tools.

The DevOps View And Cloud Migration:

DevOps teams are increasingly craving autonomy, in order to roll out applications on demand. A natural association with this approach is that of application workloads becoming more decentralised, in line with the general data/application-centric view of the IT world. However, while that itself might create operational difficulties, it's not simply a case of getting applications out as quickly and efficiently as possible but, at the same time, fully compliant and secure, the latter concerns largely being anathema to the traditional developer! Hosting them in public cloud environments adds yet another layer of potential complexity to manage, yet current deployment direction often translates into applications being deployed into a mix of public and private cloud infrastructures.

Where a company is holding back from cloud-based deployment, this might be purely due to uncertainty, not least the aforementioned security and compliance implications, and general concerns about migration. So, any solution which acts as a secure, optimal delivery mechanism is not only directly assisting with the management of a contemporary application environment, but it could act as an enabler for those companies looking to move to a hybrid architecture but are reluctant to do so.

Licensing Options

Where many companies come unstuck when scaling load balancing operations is in the licensing; all too often it's too restrictive, meaning under or over-provisioning, which is costly either way.

With its AX solution, Kemp offers a choice of perpetual, subscription or metered. You can buy per-instance or per-app, so if you know your long-term requirements and you are used to procuring hardware or virtual instances, then perpetual might be the go-to option, However, the modern DevOps approach now becoming standard practice for many companies means that metered licensing might be a better bet. For Virtual LoadMasters (VLMs), these can be licensed using permanent or subscription licenses on each instance or with metered licensing across multiple VLM instances. The latter provides the flexibility to deploy and retire load balancing resources on-demand, thereby simplifying DevOps environments and application scaling. Cloud-based LoadMasters can be purchased on a pay-as-you-go basis from the cloud marketplaces. They may also be deployed with perpetual licenses, annual subscriptions or under a metered license.

Key to metered licensing is that there is no limitation on the number of instances deployed – billing is based on aggregate peak usage. It works as follows:

- A customer takes out a metered licensing subscription (monthly or annual) at a tier ranging from as low as 1Gbit.
- The customer can then deploy metered LoadMaster instances as needed and optionally cap the performance.

- The peak service traffic (in Gbit/sec) for the month is measured for each deployed instance.
- The peaks for each deployed LoadMaster are added together to give the aggregate monthly usage.
- If the aggregate monthly usage exceeds the subscribed tier, overage is charged.

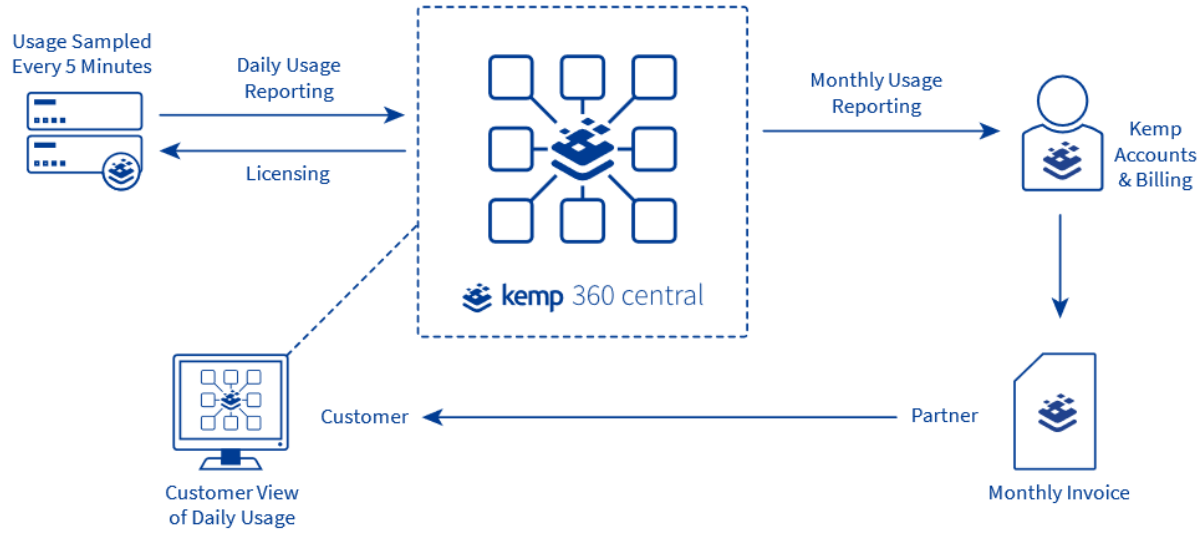

Figure 2 – How Metered Licensing Works

Cloud-Native: Partial Or Complete Solution? What About 3rd Party Appliances?

There has been much discussion about the effectiveness of cloud-native, or cloud-delivered load-balancers, such as those available from CSPs like AWS and Azure. As part of the packaged offering from each CSP, there is an undeniable convenience factor, so long as the scenario is that of a single application and a single CSP. But what might appear to be "smart cost-saving" can equally result in a lock-in to that CSP and a lack of flexibility; it's also easy to run up huge bills with the likes of an AWS as and when you need to keep adding resource and capacity – both somewhat ironic when you consider the primary rationale for moving to the cloud in the first place. Add in a second or third CSP and you are then looking at multiple, individual load balancing solutions, with multiple management headaches. And still, the features and functionality will be extremely limited.

One key feature of the Kemp AX solution, other than it being able to load-balance deployments from virtual, to hardware-based, to cloud offered in the AWS and Azure marketplaces and beyond, is its analytics support for 3rd party vendors, such as F5/NGINX, AWS ELB and HAproxy.

So, here's the ideal scenario Kemp is aiming for: an unlimited number of load balancers for any number of applications, with unlimited throughput, only paying for what you use, and continuously ensuring an always-on application experience. That's a far cry from the limits of traditional solutions in this space.

THE KEMP AX SOLUTION: HANDS-ON

Our Test Environment

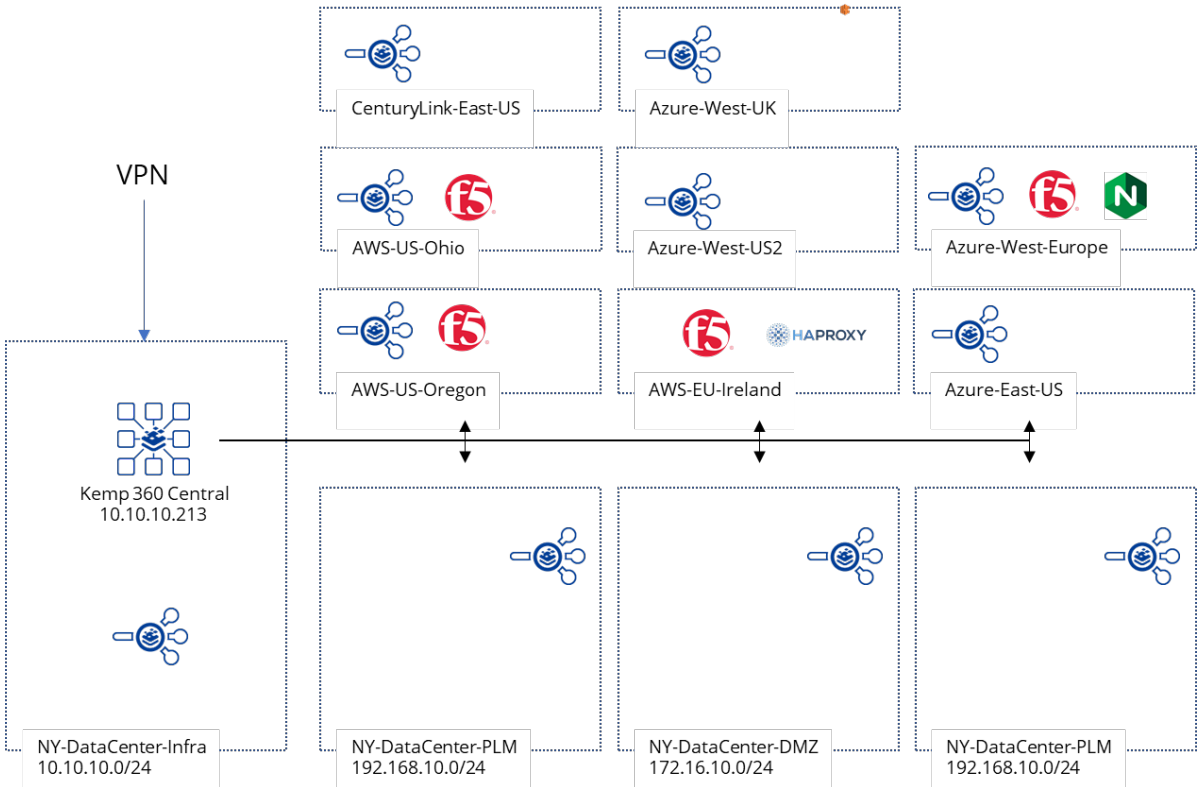

Figure 3 – Our Test Environment

In order to get some hands-on experience with the Kemp AX solution, we created access to a test environment, including a multi-vendor architecture and remote access, including VPN and RDS connectivity, from the Broadband-Testing offices.

This enabled us to validate the single point of management claims, the ability to orchestrate and manage the AX landscape from anywhere – we were managing a global, hybrid network from the UK – and in a multi-vendor/cloud environment.

AX 360 Central – Deployment And Management

Licensing/Auto-Deployment

We've already mentioned the importance in a contemporary load balancing environment of flexible licensing. The whole point of optimising application access and delivery is to remove constraints *without* incurring huge costs through over-engineering.

Within our test environment, we looked at the metered licensing option, that being especially flexible, not least because you can specify a cap or not. With any metered service, there is always the possibility of 'bill shock' where usage greatly exceeds the contracted amount and results in overage charges, similar to the old charging mechanisms for Internet/leased lines where bursty traffic was supported but charged for at mind-numbing rates. So, in order to reduce this risk factor, Kemp provides two mechanisms:

Instance Capping – this is where a customer can limit the capacity of any instance. For example, development and test load balancers may be capped to avoid accidental usage spikes, while production services may be uncapped to guarantee application availability.

95th Percentile Billing – in this case, customers with annual metered license contracts benefit from 95th percentile billing where the top 5% of usage data points are discarded, thereby removing excessive spikes from billing calculations.

Using Kemp 360 Central, we accessed a provisioned VLM instance in the NY Datacenter-PLM DC as shown above, with an instance 192.168.10.40:

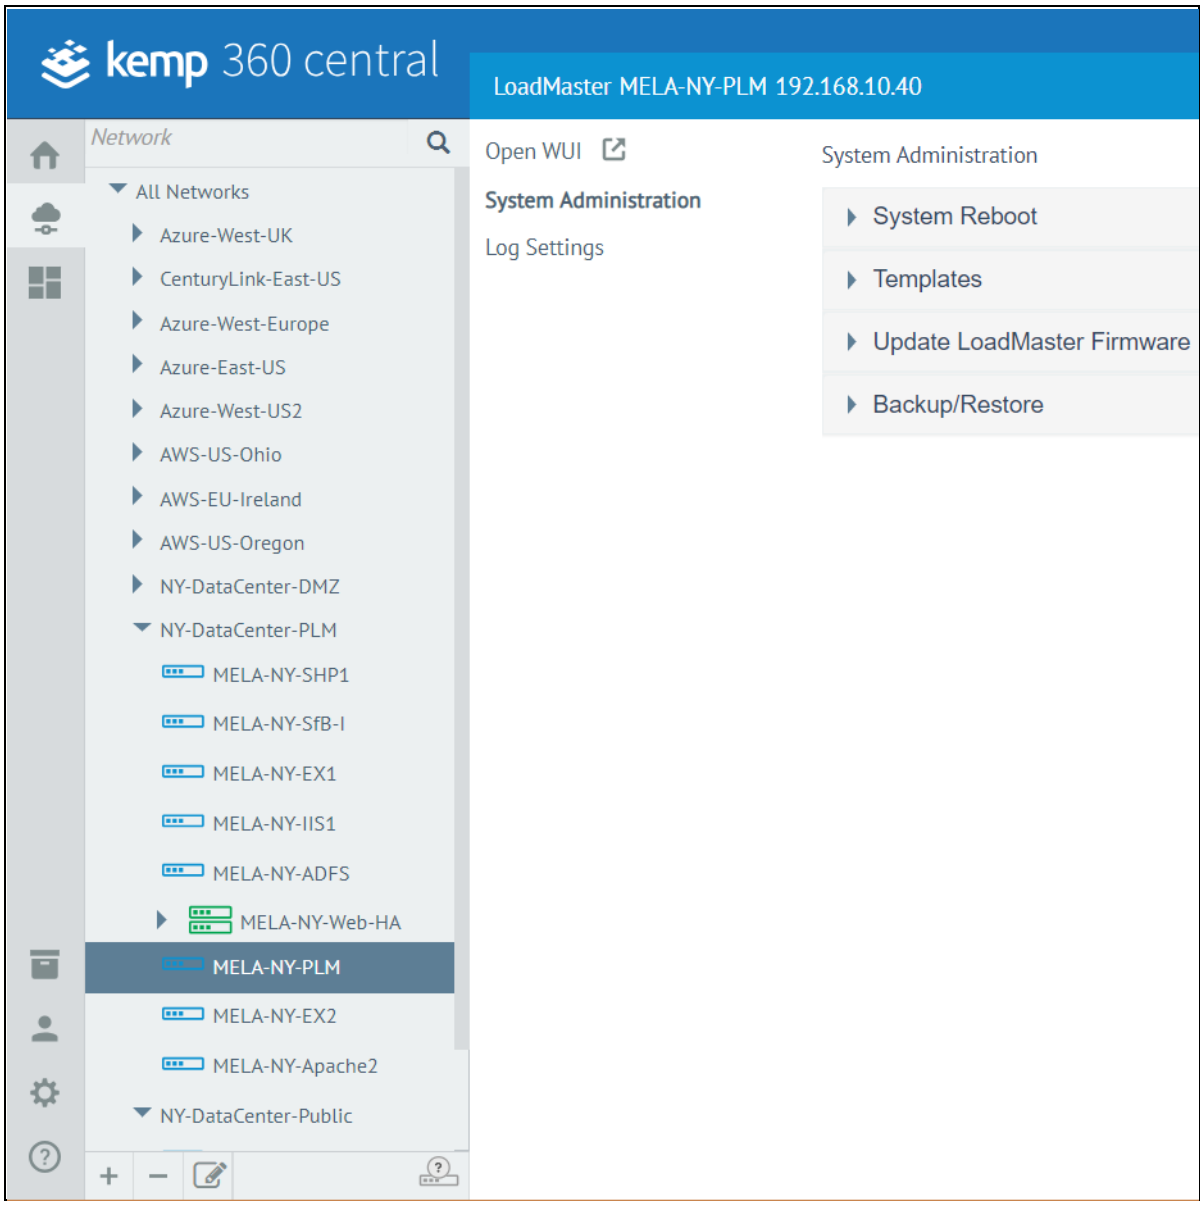

Figure 4 – Kemp 360 Central

From here we were able to access the VLM WUI and configure directly. Looking at the metered license options we had multiple choices, including uncapped. Having selected the best option, you simply complete the process of activating that license and the VLM becomes instantly available.

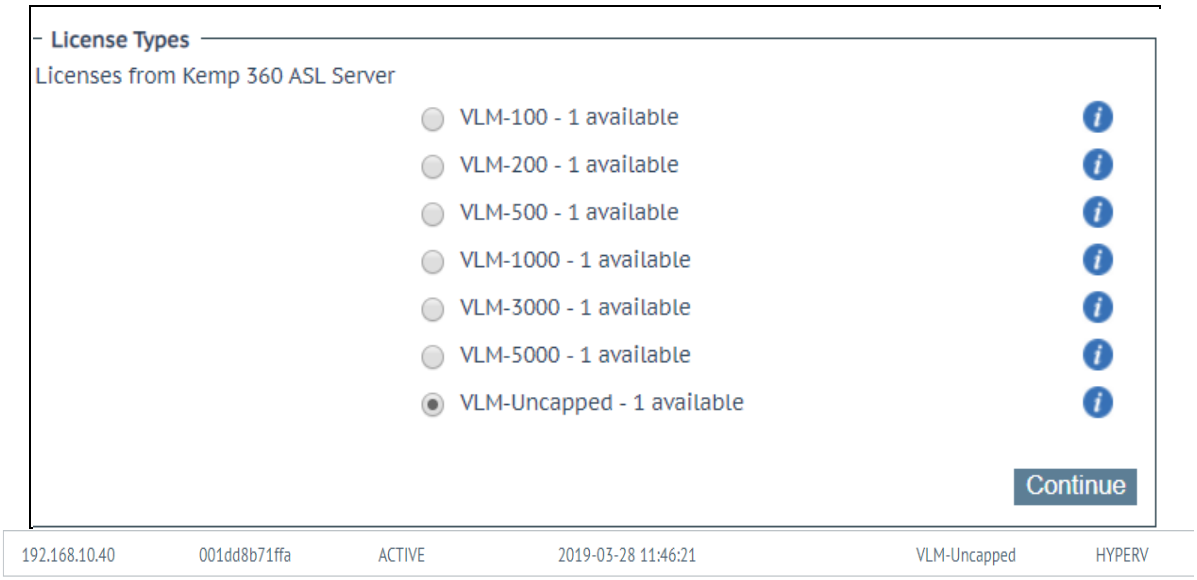

Figure 5 – 360 Metered Licensing Options And Activating A License

Importantly, license management is equally flexible. For example, unlike a traditional license pool, there is no need to recover a license for reuse. If an instance is not in use it shows as either an active or deleted license but does not incur any billing.

For deploying further LoadMasters, two basic options exist:

Quick Mode provides pre-defined application templates for common workloads and provides an easy way to get a system up and running.

Advanced Mode allows complete configuration of all LoadMaster and application settings.

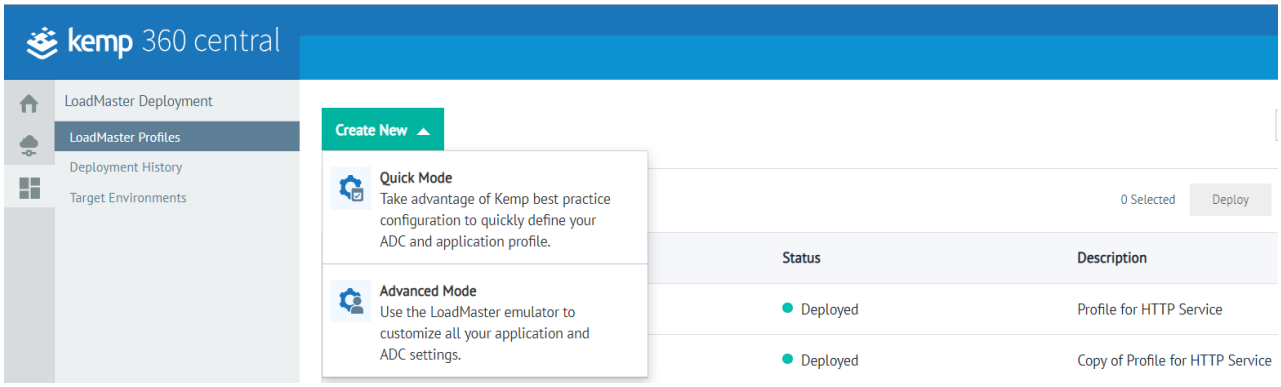

Figure 6 – LoadMaster Deployment Options

Using "Quick Mode" to deploy a LoadMaster in VMware to load balance traffic between two web servers over HTTP took literally a couple of minutes, especially when combined with automatic deployment (as opposed to a manual option). Basic information such as IP addresses (web server/VIP/gateway) plus login details and what LoadMaster mode (standalone in this case) is all that is required in Quick Mode and, for auto-deployment you have the option of creating a new 'Target Environment' or selecting from a list of previously defined targets, which is the option we chose.

Throughout the process, a progress bar at the top of the screen ensures you are completing all the necessary steps during deployment, until the deploy confirmation screen is reached:

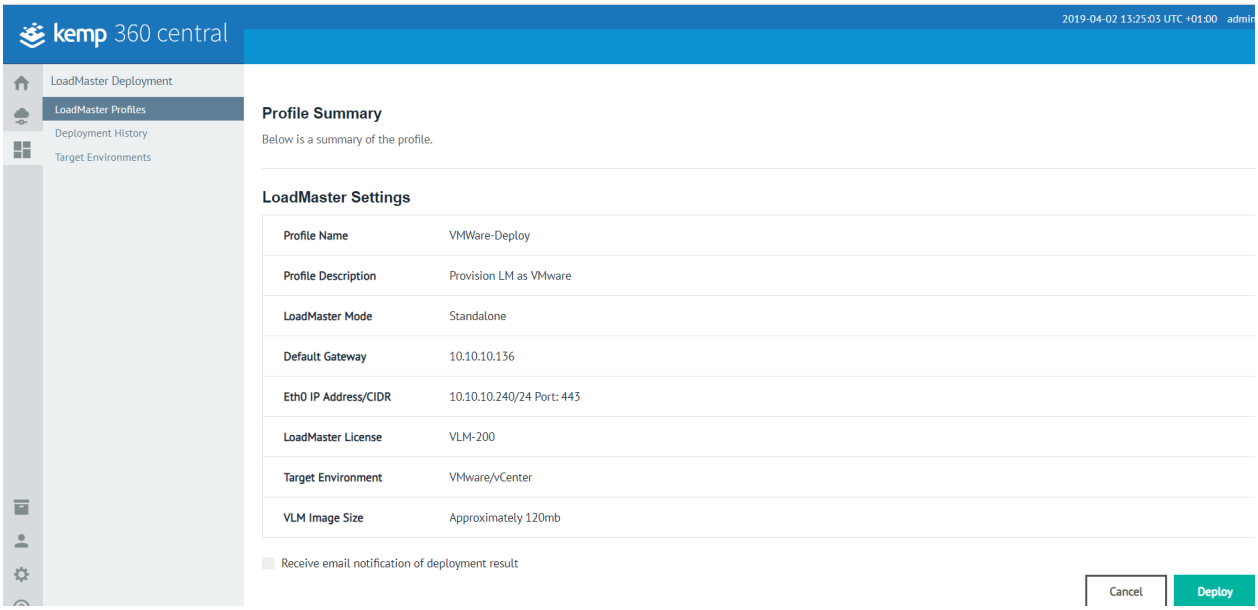

Figure 7 – Deploying A LoadMaster

Once deployed, the LoadMaster instance became active in Central, from where we were able to monitor and managed it, as with the pre-defined instance we licensed earlier. The key point to make here is that, regardless of LoadMaster type, configuration has been simplified as much as a possible and the use of templates allows for mass deployment with minimal human requirement – thereby reducing the risk of erroneous (and costly) deployment, as well as the significant time-saving and speed of deployment benefits.

Migration From OnPrem To The Cloud – VS Motion In Action

As we noted earlier, Kemp's AX solution could effectively be an enabler for companies hovering between OnPrem and cloud currently, as it simplifies the process of managing a hybrid environment.

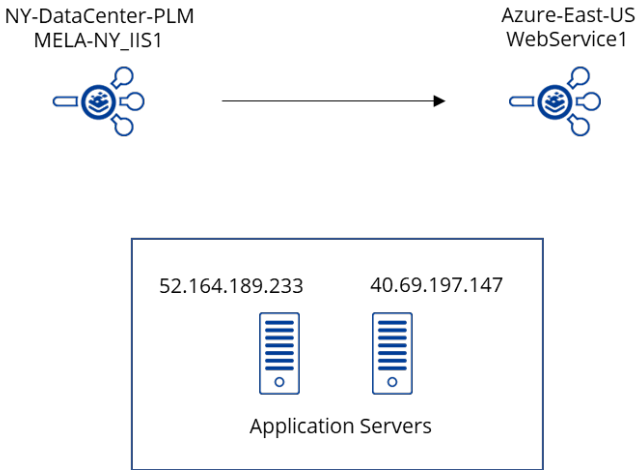

Figure 8 – Migrating From OnPrem To Cloud

For our test example, we chose to migrate a virtual service from an OnPrem appliance to a cloud-based load balancer. The VS Motion feature in Central allows a virtual service to be migrated between load balancer instances, so for this test we had an AppService on the load balancer MELA-NY-IIS1 server in our NY-DataCenter-PLM DC, with a view to migrating it to the cloud as WebService1 in the Azure-East-US cloud shown in our test environment diagram and above.

Until recently, this kind of migration would be seen as – frankly – unlikely! However, the ability to treat the entire environment as one giant load balancer landscape makes this kind of task trivial and took just a handful of simple administrative steps – note: nothing technical required. We simply selected this source, went to the 'Service Configuration" tab on Central, selected the "VS Motion Migrate" icon for the WebService1 virtual service, set it as the target LoadMaster, specified an IP address, clicked on the "Enable VS on target" box and clicked on "Move".

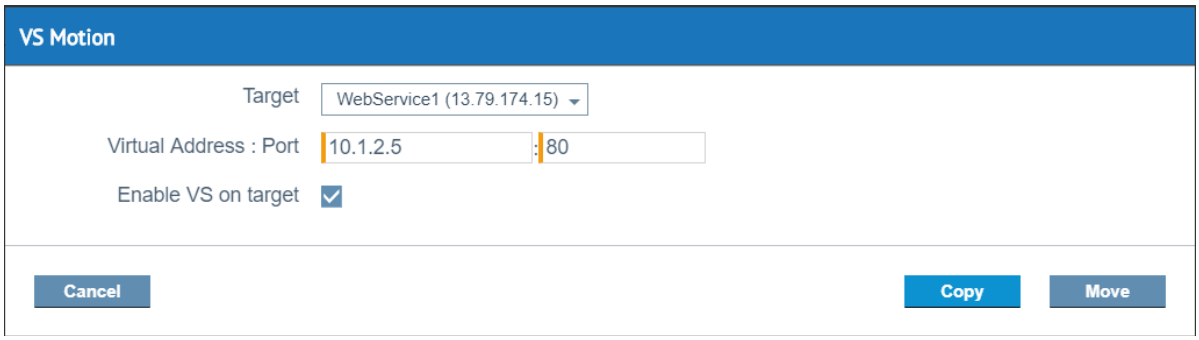

Figure 9 – Migrating Webservice Using VS Motion

We also did a reverse migration, from an Azure cloud (MELA-AZ-HR from the Azure-West-US2 cloud) to the NY Datacenter-Infra (MELA-NY-HR). Key here is that the NY Datacenter we migrated to has two interfaces – a virtual and a real, but we moved the cloud-based service to the real server interface, showing the full flexibility of the Kemp solution.

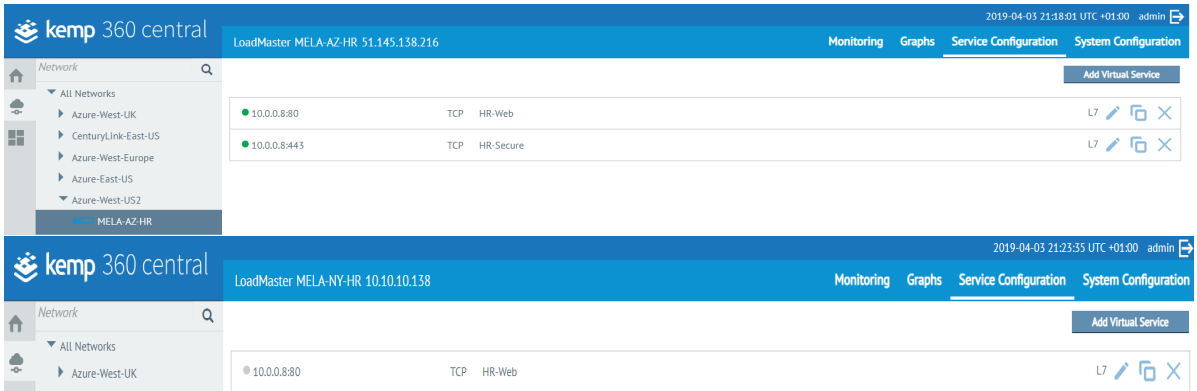

Figure 10 – Migrating Cloud Webservice To Real Server In NY Datacenter

Monitoring with Kemp 360 Central

In addition to being a deployment, management and orchestration tool, Kemp 360 Central has further primary uses for monitoring the status of the load balancer estate. Foremost is that global dashboard, which gives a top-down view, indicating the overall health and performance of load balancer assets and highlights instances with the highest CPU usage, for example. It gives you the ability to dig deeper via a number of onboard

windows such as Device Health, Local Licenses, Non-Local Licenses & Subscriptions, Top 5 Utilization and Log Summary.

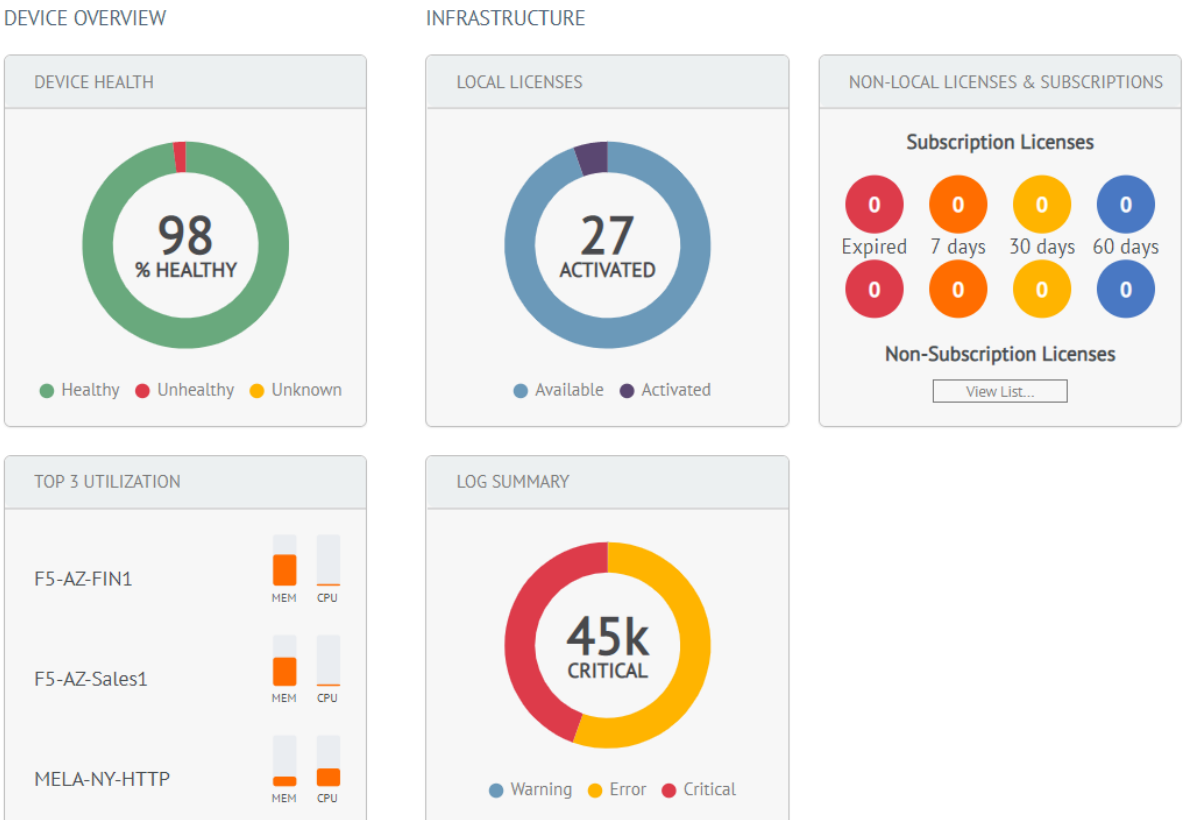

Figure 11 – Part Of The 360 Central Dashboard

Hovering over any sub-window proves more information, while click on that window, such as Device Health, drills down into more detail:

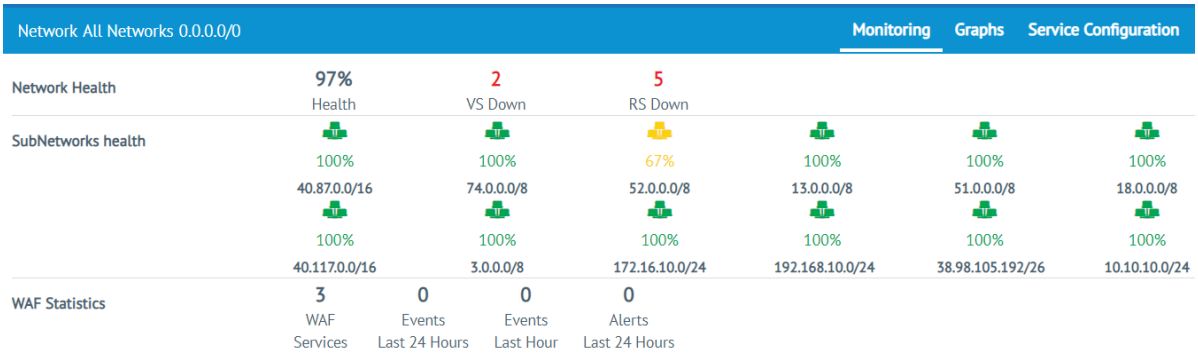

Figure 12 – Drilling Down Into Device/Network Health

Additionally, there are Health infographics, including Real/Virtual Server Status, WAF Statistics and Active Connections, to provide a real-time view of Virtual Service and Real Server status and their connections, as well as any security issues.

Multi-Cloud, Multi-Vendor Management

As we noted earlier, Kemp 360 Central provides multi-cloud management and control with support for On-Prem hypervisors, AWS, Azure and CenturyLink clouds. For example, in our AWS-US Oregon DC, we were monitoring an AWS environment with F5 BIG-IP appliances active.

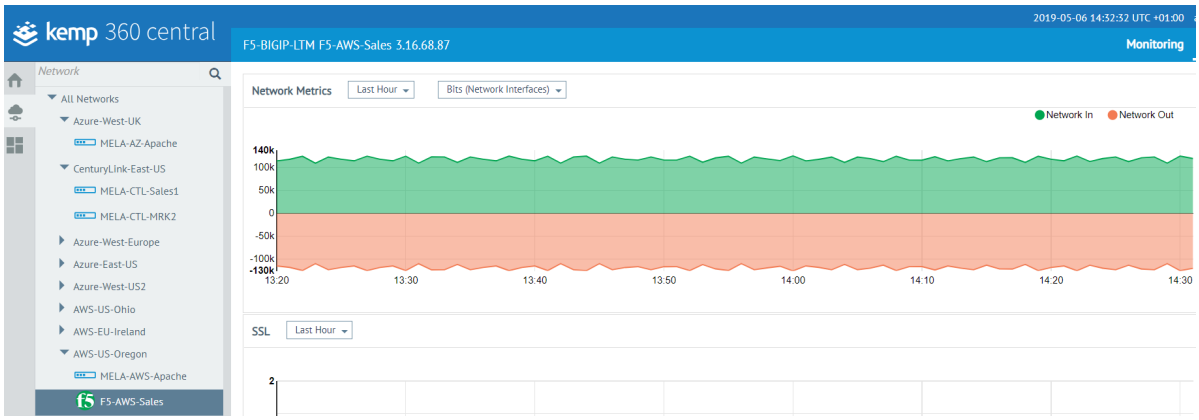

Figure 13 – AWS And F5 Appliance Monitoring

In addition, on our test network, the AWS-US-Ohio and AWS-EU-Ireland clouds hosted F5 and HAProxy ADC instances, as well as an ELB (Amazon's Elastic Load Balancer) instance configured at AWS-US-Ohio.

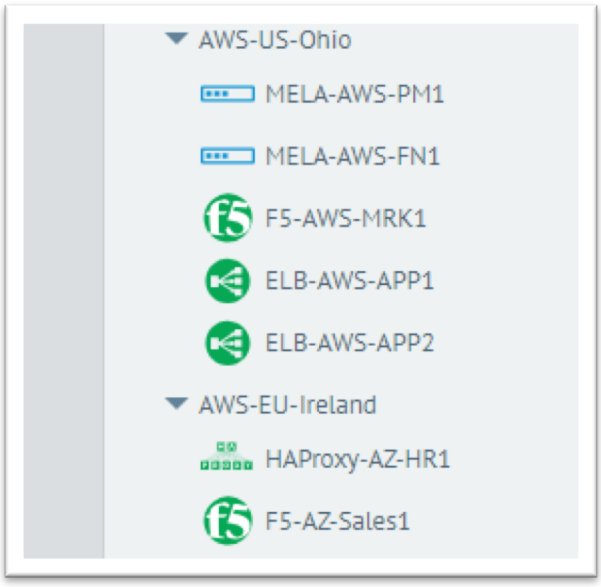

Figure 14 – Global Multi-Cloud, Multi-Vendor Monitoring

Kemp 360 Central And Vision – Troubleshooting Problems

Kemp 360 Vision

As the next level of analytics after Central, Kemp 360 Vision does more than simply flag up problems, but also suggests possible remediation action and also feeds into additional communications such as the Kemp Slack channel we used extensively during the testing. As an "in your face" instant red flag however, it does this job perfectly, noting here two problems that we generated for testing purposes – see later in this section for more details!

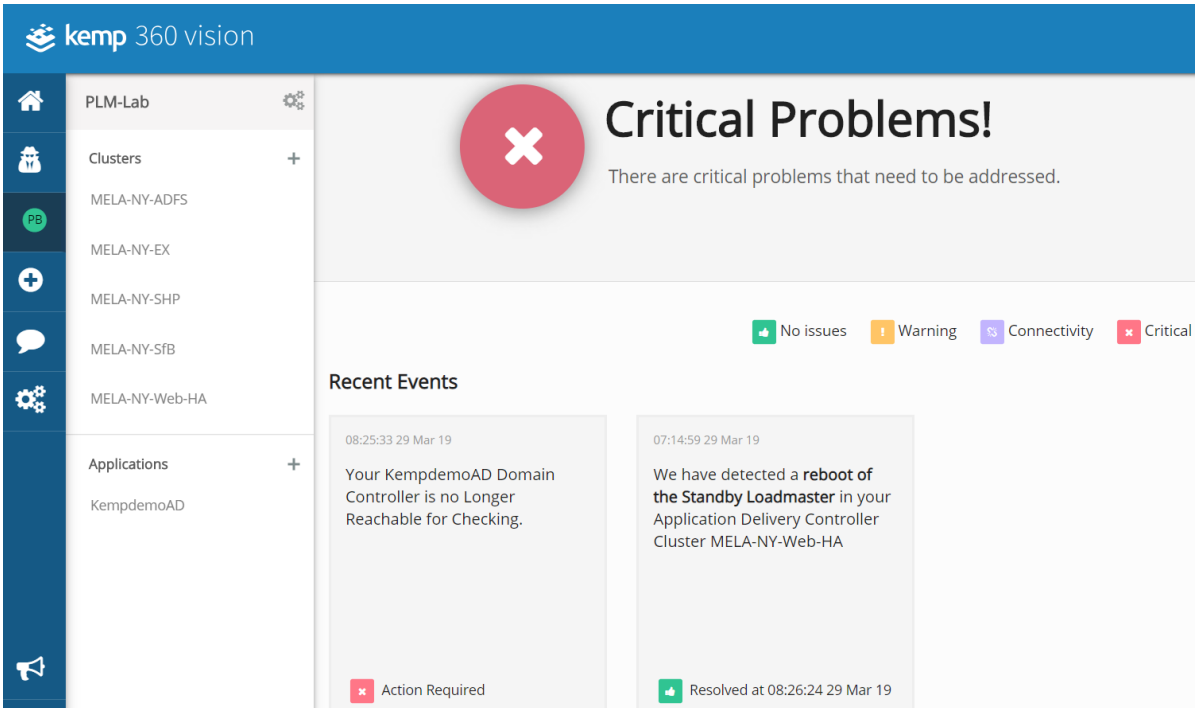

Figure 15 –Vision in Action

Equally, it advises once a problem is resolved and provides a continuous update of the overall status of the load balancing estate:

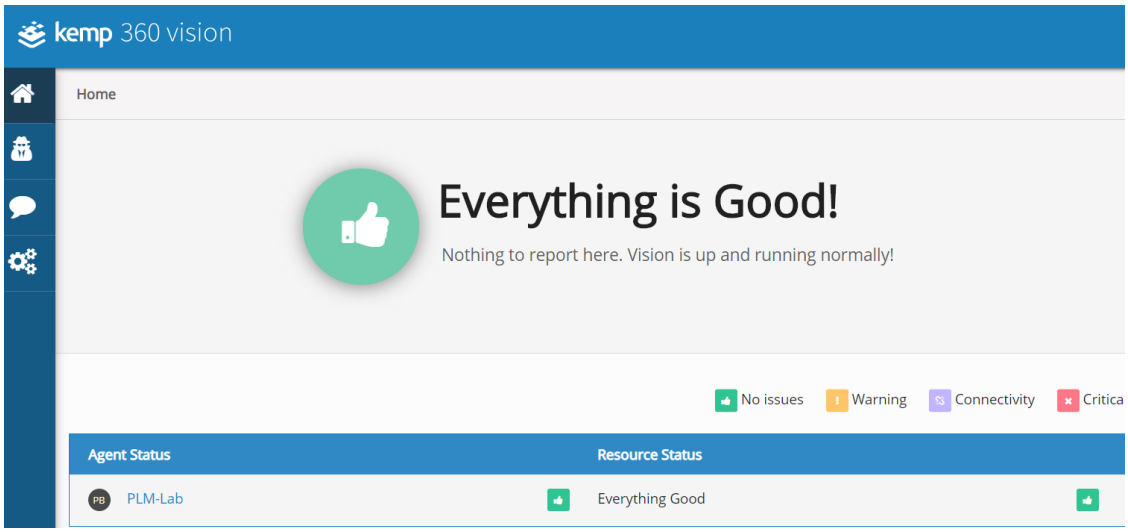

Figure 16 – Kemp 360 Vision Updating Status Positive

Server Overload Identification

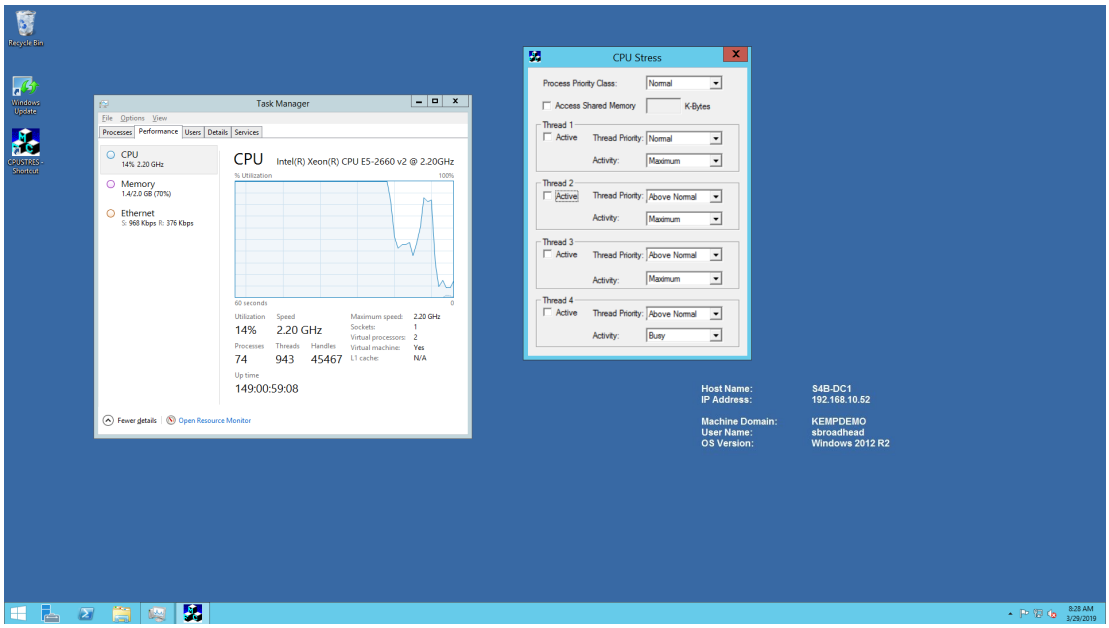

Figure 17 – Creating An Overload On An AD Server

We'll now look at how both Kemp 360 Central and Kemp 360 Vision – and the AX comms channels - can be used to analyse problems, uncover and help resolve issues, starting with a classic problem – a server overload. In this scenario, we remoted into a Windows Active Directory server and created traffic using a public domain application stress tool in order to saturate the CPU load and generate an issue – see above graphic.

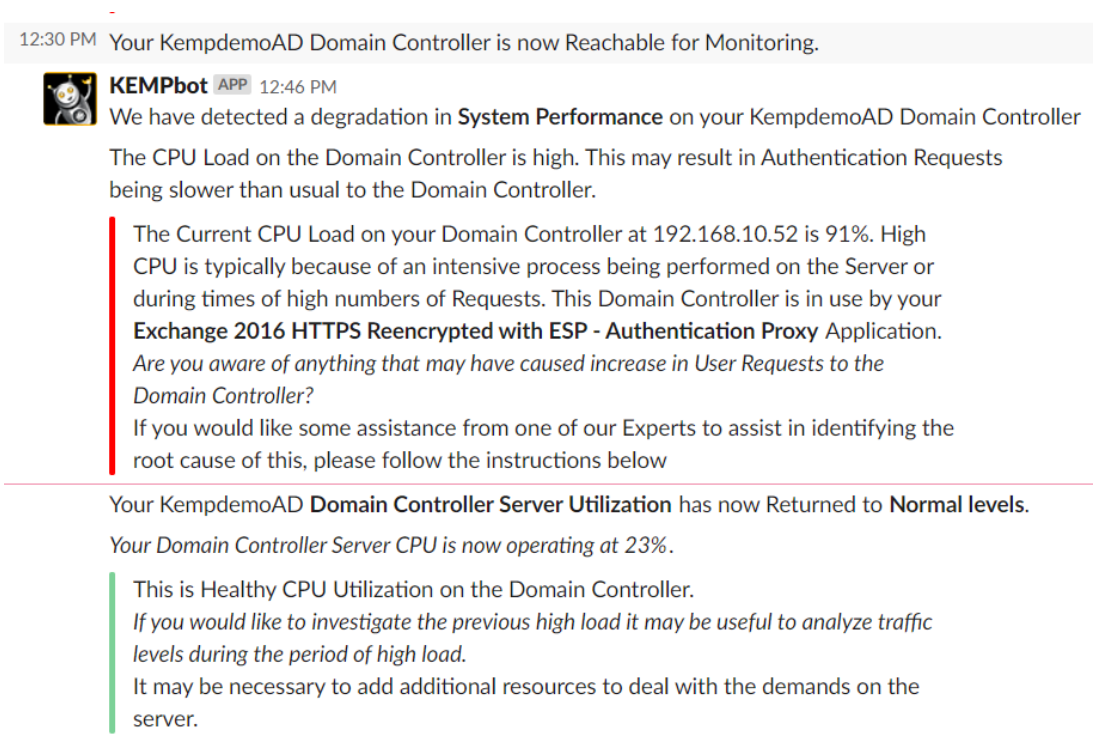

Figure 18 – Kemp Slack Channel Generating Alerts And Updates

Having let the stress tool run for several minutes, we then stopped the test. Meantime, we received an alert via our Kemp Slack communications channel, indicating an issue with our stressed server, showing the CPU utilisation rates (91%) and suggesting reasons for the high utilisation, noting that user requests are especially high. Once the AD server utilisation rates returned to normal, we received a follow-up message via the Slack channel (see above) indicating a healthy server condition.

An alert will appear on Slack indicating the resource usage status. As can be seen in the example alert below, not only is the high usage alerted, context is also provided on the applications that are using the AD services.

Kemp 360 Central/Vision: HA Failover

This time our focus was on a combination of Kemp 360 Central and Kemp 360 Vision, when we initiated a High-Availability (HA) failover situation.

We did this my initially connecting to Central, navigating to the NY-DataCenter-PLM network and select MELA-NY-Web-HA. The monitoring screen for this HA pair indicated which instance was Active and which was Standby:

Figure 19 – Kemp 360 Central: HA Pair Status

From the network tree, we selected the Active member of the HA pair and issued a system reboot command, to immediately reboot the instance. This created an HA failover, which was recorded via the Kemp Slack comms channel.

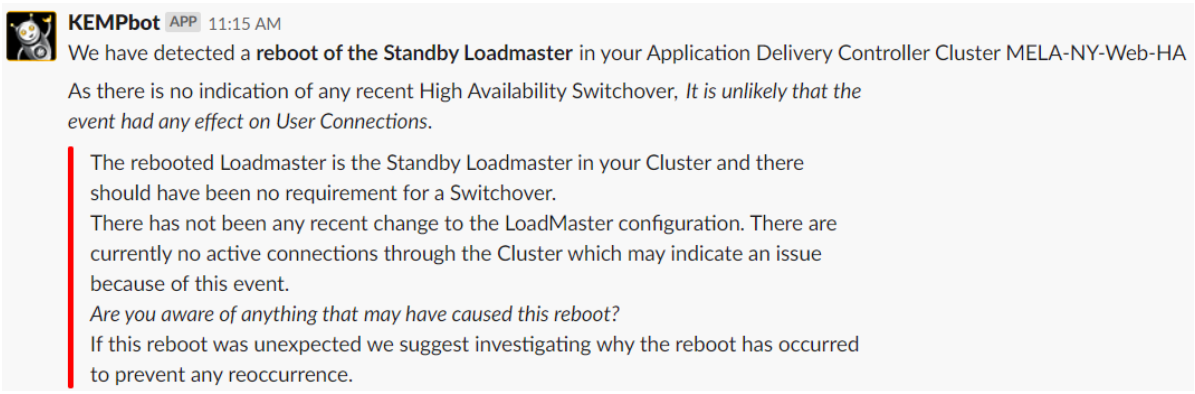

Figure 20 – Kemp Slack Channel Identifying System Reboot

As part of that alert, Vision had fed in additional contextual information, notably that that there had been no recent configuration changes and that there were no active connections. Obviously, in this case, as we've engineered the failover, we know the causes, but this would otherwise have been valuable information, indicating a potential hardware issue, for example.

This information was obviously also available directly from Vision itself:

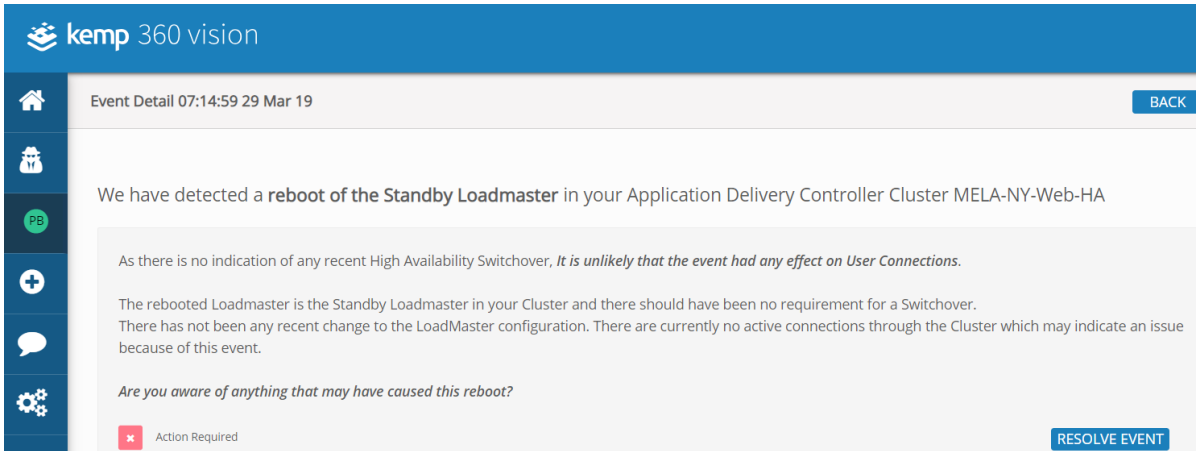

Figure 21 – Kemp 360 Vision Noting Reboot And Feeding Into Slack Comms

Cross-checking with Central, we could see that the HA failover worked immediately, as it should with the previously standby instance, now active:

Figure 22 – HA Standby Now Active

What we saw here is how AX combines various elements – in this case Central and Vision and the Kemp Slack comms channel – to work as a unified troubleshooting and management toolkit.

Pro-Active Identification Of Application Problems

Not all "problems" on the load balancing estate are immediately critical but can be slow-burners that need identifying and resolving before they can become critical.

One example we looked during testing was uneven load-balancing, where we deliberately engineered a situation where traffic was being wrongly directed to one server and used AX to identify this. We created a user case with a LoadMaster - MELA-NY-HTTP - delivering traffic to two servers based on cookies, where the source traffic was coming from the same IP address. By changing the LoadMaster configuration to use source IP address scheduling, all traffic would end up on the same server and trigger an alert.

We could see, graphically, the build-up in traffic on the one server; simulating an investigation it was immediately obvious that the resolution was to change the load balancing algorithm, in this case to a round-robin method – see over page.

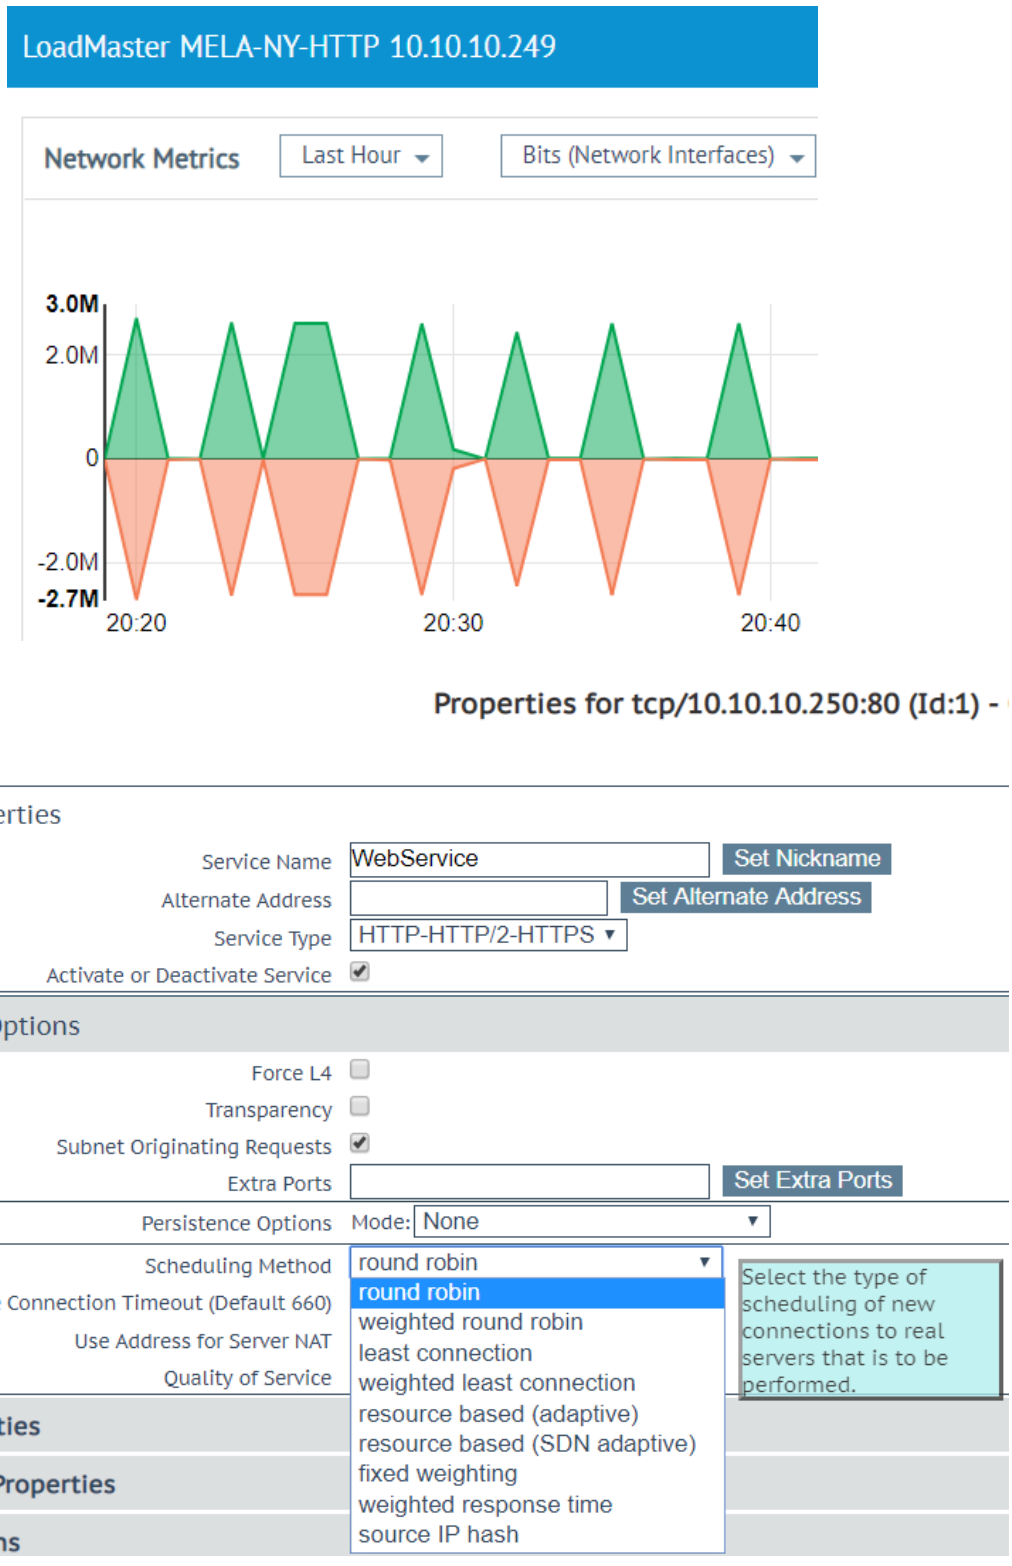

Figure 23 – Uneven Load-Balancing: Problem And Resolution

Another very common problem experienced by IT support teams, usually as a result of a call from an irritated user, is that of slow application response times, typically incurring a long round trip response time (RTT) so, for the last element of the testing we created such a scenario to see what the AX solution could tell us about such a situation.

In this instance we have noted that, of three servers providing application delivery, one is performing particularly badly. RTTs have shot up, so those users being served by that particular server are experienced very poor service, while others are not experiencing any problems.

This – in the real world – is a very difficult problem to identify; it might be a problem at the client itself, the client to load balancer connection, on the load balancer itself or the connections from the L-B to the real server. That's a lot of potential problem areas to cross-check. And since, in our test example, the real server is still actually running, there are no obvious alerts to spot.

However, using LoadMaster analytics and Vision, we were able to see exactly what the problem was and where it was occurring. First, we note especially high CPU activity on one of the servers, using the LoadMaster interface.

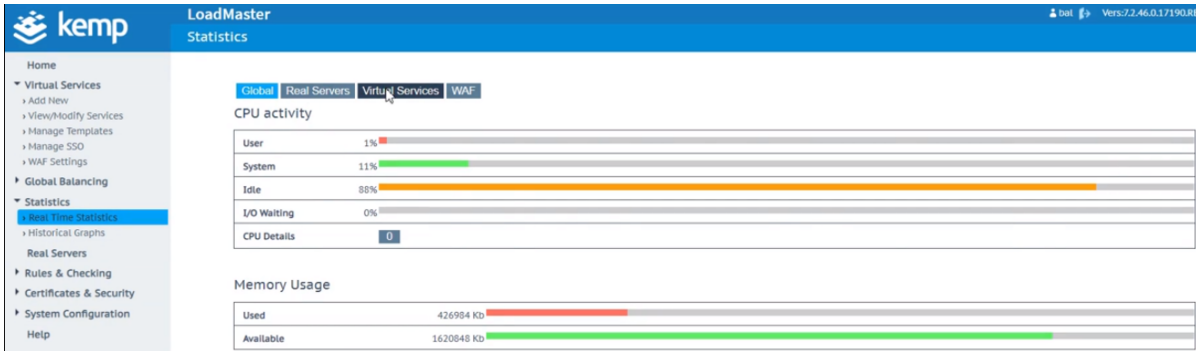

Figure 24 – High CPU Activity On A LoadMaster

At the same time, we received an alert from Vision, identifying that there is a problem.

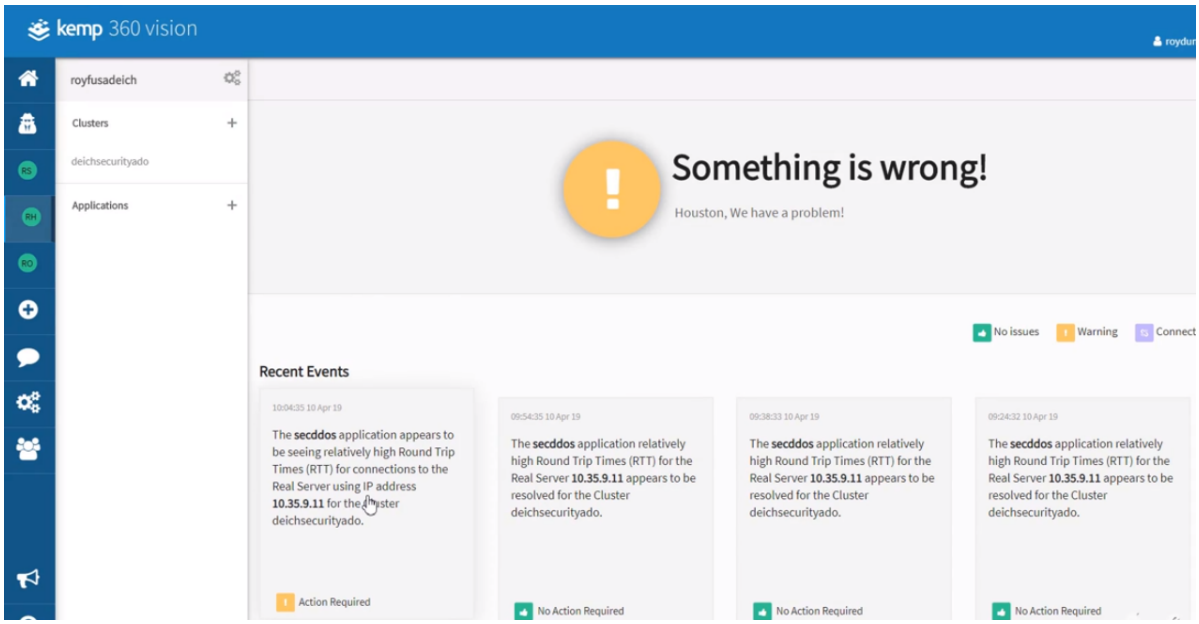

Figure 25 – Kemp 360 Vision Sends Alert

We were now able to drill down to carry out more analysis on the problem area that Vision identified,

These included previous occurrences – and we could see how they had suddenly started rising - – a clear indication of a potential server problem. Note: this alert and related information, as in previous examples, could also be passed on from Vision by other communications channels, from Slack to SMS.

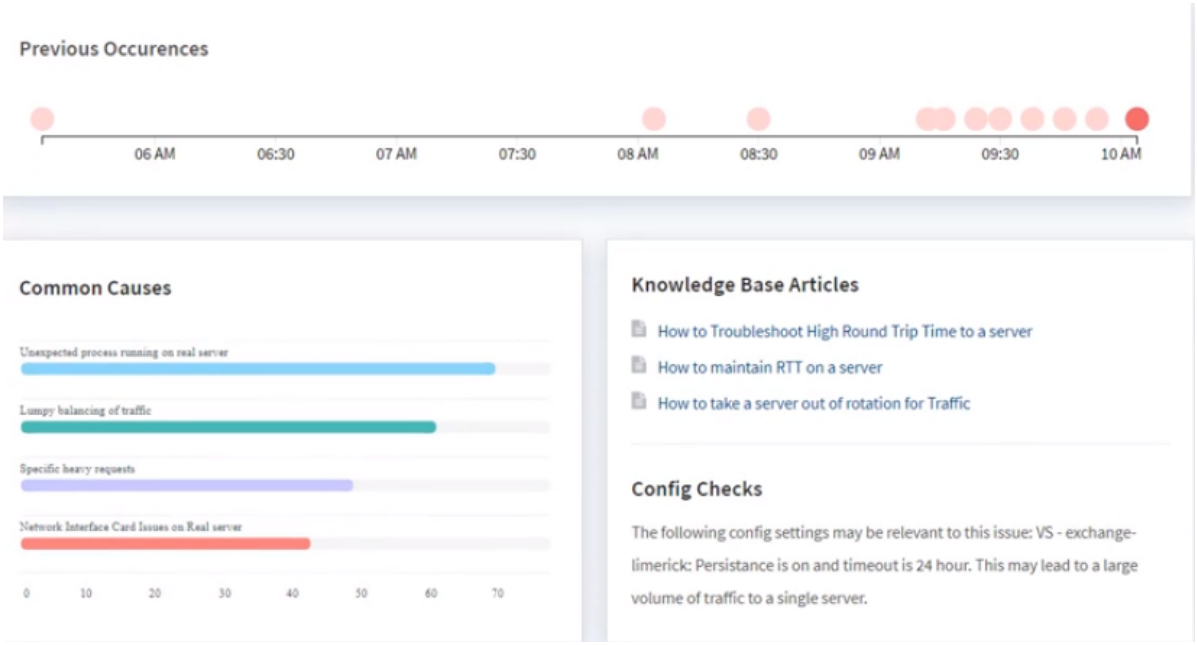

Figure 26 – Identifying Problems And Possible Remediation

Vision also suggested remediation possibilities, by linking to the Kemp knowledge base library, as well as identifying common causes.

IN CONCLUSION

Given the changes in contemporary IT architectures, traditional load balancing architectures simply cannot be expected to deliver any longer, and Kemp clearly understands this.

Its AX fabric is a clear indication of a new direction in application delivery, supporting OnPrem, cloud and 3rd party vendors within a single solution and the ability to drill down into the minutiae from a top-down perspective. Similarly, covering off optimisation, analytics and security is an equally important factor to consider.

The world is moving from a data-centric view of IT to an application-based approach – something the new generation of users has all but insisted upon – and the DevOps up approach is equally prevalent. So managing and optimising IT is no longer about the technology but about supporting everything, everywhere as one giant piece of IT real estate. And from a single console that, itself, can be located anywhere.

Just as important is the ability to control and save costs for the customer, especially in light of the hidden costs of cloud adoption and agile DevOps. Again, Kemp seems to have grasped this concept and AX is a huge step in the right direction – definitely a "must take into consideration" platform for anyone looking to update, expand, or simply get into the world of load balancing.