USER GUIDE

V3.3

UILA INC

uila

Table of Contents

1. Introduction

1.1. Scope and Purpose

The first part of this document describes the system requirements, installation and configuration steps for Uila software.

The second part details how to use the console in order to manage and troubleshoot application and infrastructure related issues in the data center.

It is assumed reader is already familiar and proficient in VMware installation, configuration and on-going management.

1.2. Architecture Overview

Uila consists of three major components –

- *Management and Analytics system(UMAS)* – UMAS is a big data store and analytics engine that was designed to accommodate a large data center deployment with thousands of servers. UMAS can store data for up to 1 year and record data in minute resolution, while maintaining real time responsiveness. UMAS's built-in redundancy offers high availability, removes downtime and reduces maintenance overhead.

  The UMAS can be used to provide a single pane of glass view for end-to-end visibility into performance, capacity and resource usage/allotments across all on-premise and cloud-hosted services (VMware, Microsoft Hyper-V, Amazon Web Services, Microsoft Azure, Google Cloud and other cloud provider deployments).

- *Virtual Information Controller(vIC)* - vIC is installed as a virtual appliance on-premise or on cloud datacenter. The vIC retrieves the infrastructure configuration and collects network, storage and compute performance metrics. This is then combined with the data from the vST and transmitted to the UMAS.

- *Virtual Smart Tap(vST)* – vST is deployed at the host as a small foot print guest VM that utilizes Deep Packet Inspection (DPI) technology to identify unique applications and its attributes. The vST measures application response time and collects network performance data. No packet payload is examined or stored, thus removing the risk of exposing sensitive data.

  In a cloud deployment, the VST, also collects the network and performance metrics from the IST and utilizes the Deep Packet inspection technology to identify applications.

- *Instance Smart Tap (IST)* – The Uila Instance Smart Tap (iST) is deployed as a plug-in in a distributed manner across the Public Cloud on the VMs or Instances running the application workload. It collects traffic as well as VM and Instance level Compute statistics and sends it to the vST for Deep Packet Inspection.

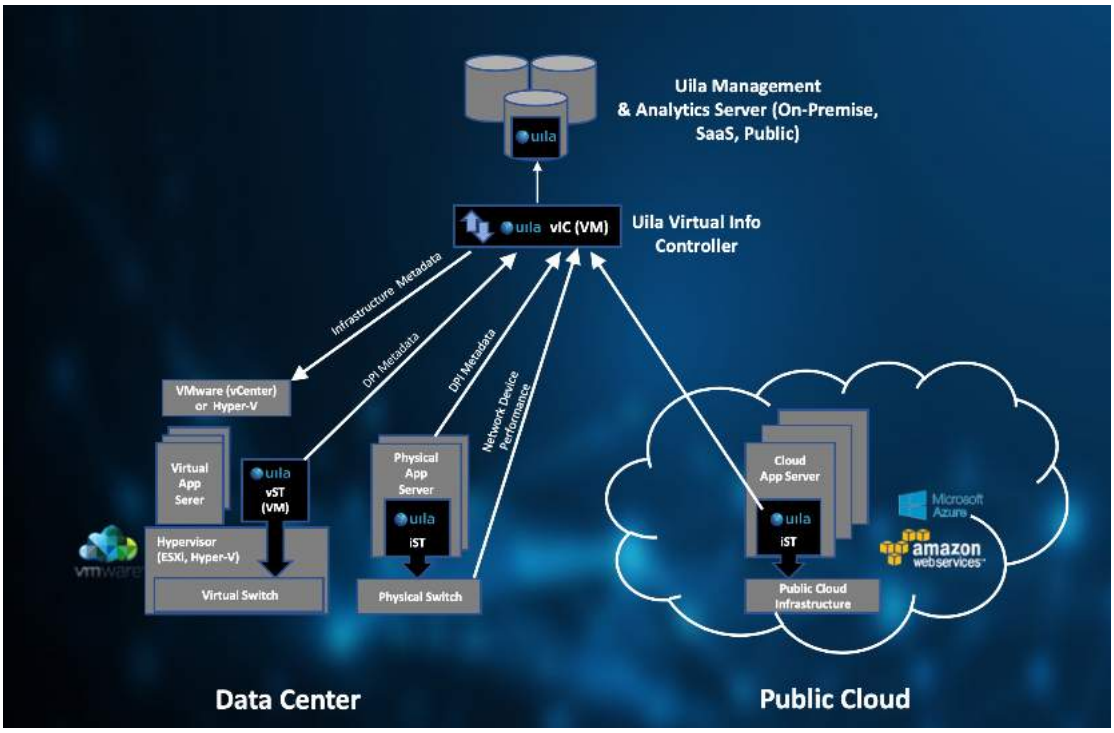

Fig 1.1 – Uila Architecture overview for Multi-Cloud

1.2.1. Unified Central Management Console

Modern virtual technology has improved data center's operating efficiency. However, the management tools that IT organizations use may not effectively cope with the increase in complexity to monitor application performance. Uila management console dashboard offers a simple yet powerful view to visualize the heath of an Applications across a Multi-Cloud environment. It also reveals the underlying physical/virtual infrastructure in the network, compute and storage segments to pinpoint the application performance degradations and bottlenecks.

1.2.2. Automation and Provisioning

To aid data center operators, Uila integrates closely with the VMware vCenter and cloud platforms such as Amazon Web Services, Microsoft Azure, Google Cloud, VMware Cloud on AWS, Alibaba Cloud to setup applications and tenants for monitoring. Uila can also configure, deploy and provision the Uila guest VM's automatically, that eases the additional burden of maintenance and support.

1.3. Feature Highlights

1.3.1. Multi-Cloud Architecture – Built for Multi-Cloud Data Center

Uila architecture is a next-gen platform that utilizes the latest big data technology which offers unprecedented scalability and flexibility to monitor mission critical business applications across the multi-cloud cloud, while maintaining real time responsiveness:

- Scales from small to large data centers with built in redundancy for high availability.
- Maintains historical records of up to one year.
- Small footprint virtual Smart Tap(vST) with minimal overhead is deployed as a guest VM for on-premise datacenter.

- Low resource utilization Instance Smart Tap(iST) with minimal overhead is installed into a VM/Instance for the cloud datacenter.
- Collects application response times with more than fifty critical infrastructure performance metrics in minute intervals.
- Embedded Deep Packet Inspection (DPI) technology to identify over 3,000 unique applications and their attributes.
- The vIC seamlessly integrates with the VMware vCenter leveraging the network, storage and compute performance metrics maintained by it.
- Uila only collects metadata. Packet payload is not examined or stored. Data is transmitted through an encrypted SSL channel, removing the risk of exposing sensitive data.

1.3.2. Unified View – Simplify Data Center Operations

The complexity of Datacenter infrastructure hierarchy that comes with today's Multi-Cloud datacenters require an easy but powerful tool set. Uila helps data center operators visualize and pin point areas of performance degradation that can identify the root cause immediately:

- Customizable Application and Infrastructure health dashboards that mirror the logical constructs of a data center.
- Uila aggregates data into meaningful Key Performance Indicators for early symptoms of poor performance.
- Powerful analytical tool sets for Application Topology, Flow Analyzer, CPU Usage, Memory Usage, and Storage Usage provide unique diagrams that reveal the underlying impact of application performance on the physical and virtual infrastructure.
- Innovative web-based UI design which simplifies navigation and speeds up problem resolution.
- New adaptive baseline technique to enable monitoring thresholds that align with actual average performance characteristics for the underlying infrastructure. This baseline technique reduces false positives and provides accurate root cause analysis.
- Integrated alerting and troubleshooting scenario for Help Desk or Network Operation Center.
- Built-in and customizable C level reporting for service level agreement compliance.
- Exportable historical trending data as a template for future planning.

1.3.3. SaaS Cloud - Automation and Provisioning

Wide adoption of virtualization and cloud technologies have made SaaS a widely acceptable consideration for IT. As enterprise and service providers continue to seek better service and lower the cost to service their customers, Uila Cloud helps to reduce IT Operational and Capital Expenditure:

- Single pane of glass view for the performance of the Multi-Cloud.

- Integrating closely with VMware vCenter allows data center operators to take advantage of their infrastructure configuration and setup a vApp monitoring profile.
- Automated deployment and provisioning of Uila guest VM to frees up the burden of maintenance and support.
- SaaS deployment model eliminates the requirement to procure, deploy and maintain appliance and/or hardware probes.
- Multi-tenancy offers easy and common access for IT team

2. Terminology Used

This section lists common terminology used throughout the product User Guide. Uila's goal is to use the same terminology as commonly used and defined within the virtualization industry.

Terminology or Legend	Definition
Application Response Time	Time measured on the server from the arrival of a client request to the transmission of a server response.
Application Service	**Refer to Classifier
Classifier	Often used interchangeably with Application service, classifier defines the application name as a result of Deep Packet Inspection by the vST software agent. i.e. - MySQL, iMap.
Cluster	Collection of hosts and associated virtual machines. Physical resources from all the hosts in a cluster are jointly owned by the cluster and centrally managed. i.e. - vCenter Server manages the clusters in a VMware implementation.
DPI	Deep Packet Inspection uses advanced method of pattern matching and session heuristics to identify applications and their associated attributes. This helps IT organizations track mission critical applications and transaction performance issues.
DvSwitch	DvSwitch's or Distributed Virtual Switch's simplify the management of hosts in a cluster by creating a single switch across the cluster to efficiently manage multiple virtual port or dvPorts. i.e. – A single dvSwitch can apply configurations to all applicable ESX or ESXi hosts, while vSwitch can only apply configurations to one host at a time.
DvPortGroup	DvPortGroup represents a group of dvPorts that share the same configuration template. The configuration is inherited from the dvPortgroup to the dvPorts.
Host	A physical server that supports a version of hypervisor. i.e. - VMware ESXi, Microsoft Virtual Server.
pCPU	A PCPU refers to a physical hardware execution context. This can be a physical CPU core if hyperthreading is unavailable or disabled, or a logical CPU (LCPU or SMT thread) if hyperthreading is enabled. For example, a server equipped with a CPU with 4 cores without hyperthreading will have 4 pCPU. If hyperthreading has been enabled then a pCPU would constitute a logical CPU. This is because hyperthreading enables a single processor core to act like two

	processors i.e. logical processors. i.e. - if an ESX 8-core server has hyper-threading enabled it would have 16 threads that appear as 16 logical processors and that would constitute 16 pCPUs.
Port Group	It is a group of ports on a vSwitch. A 'PortGroup' is created in a Standard switch and Distributed switch. It acts as a logical segmentation of a vSwitch.
RTT	It is the time delay imposed by the networking infrastructure for a Client to get a response from the Server. The value is an average of all the TCP connections that is made to the Server.
TCP Fatal Retry	Refers to the count of retry attempts made by either the Client or the Server when it does not receive a response in a TCP conversation. A retry attempt of greater than 3 seconds and over 3 attempts is counted as a single Fatal Retry for a single minute. It is not counted again within that minute. Uila displays the count as a total, not averaged for all flows.
Tenant	Tenants can be used to provide isolation between independent groups in shared cloud environment, where multiple companies, divisions or independent groups are using a common infrastructure fabric. Tenants are useful for isolating the users, resources and services from one tenant from those of other tenants.
ToR Switch	A Top of the Rack or (ToR) switch is a high port count switch, typically 48 1G or 10G ports plus 4 additional up link ports that sits on the top of server rack in Data Centers or Co-location facilities. ToR switches are then connected to the next level aggregation switch or core router to allow communication between servers in different rack or to internet.
vApp	vApp is a collection of pre-configured virtual machines (VMs) that combine applications with the operating systems that they require. VApp's allow disparate VMs to work together in a stack as an application, and support cloud computing architectures. vApp is a VMware defined term and may be used in other similar products.
vCPU	A vCPU stands for Virtual Central Processing Unit. One or more vCPUs are assigned to every Virtual Machine (VM) within a cloud environment. Each vCPU is seen as a single physical CPU core by the VM's operating system. If the host machine has multiple CPU cores at its disposal, then the vCPU is actually made up of a number of time slots across all of the available cores, thereby allowing multiple VMs to be hosted on a smaller number of physical cores.
VM/Instance	A virtual machine (VM) or an Instance is a software, emulating a complete system platform (i.e.- a server) that supports the execution of a complete operating system (OS).
vIC	Virtual Information Manager is a Uila software agent that is implemented as a guest (VM). The vIC (1) interfaces to vCenter to retrieve compute and storage performance data, (2) acts as a proxy for vST to transfer vST meta data to Uila Cloud, (3) receives Uila management commands to install and configure vST. There is only one instance of vIC per vCenter.
vST	Virtual Smart Tap is a Uila software agent implement as a guest (VM) resides in the same Host as other application VM. It captures and analyzes all traffic between VM's within the same host, and other hosts.

vSwitch	vSwitch is short for Virtual Switch and represents networking entities connecting Virtual Machines in a virtual network at layer 2. The Virtual Switch is fully virtualized and connected to a NIC (Network Interface Card) inside a server. The vSwitch merges physical switches into a single logical switch. This helps to increase bandwidth and create an active mesh between server and switches. The VMware Virtual Switch is a switching fabric built into the VMware infrastructure (ESX) that allows you to network your Virtual Machines (VMs).
VPC	A virtual private cloud (VPC) is an on-demand configurable pool of shared computing resources allocated within a public cloud environment, providing a certain level of isolation between the different organizations (denoted as users hereafter) using the resources.

Table 2.1: Uila Terminology Definitions

3. Icon Definitions

This section lists used throughout the product and the documentation.

Icon	Definition	Usage
	Maximize display viewing area by hiding browser menu and other title bars. Toggle to restore original display view.	
	Logout your Uila session.	
	Launch help.	
	Select color for the title bar.	
	Collapse or minimize the individual sub-view within the Dashboard.	
	Restore the minimized the sub-view within the Dashboard.	
	Toggle between full screen and normal mode.	
	Re-layout the Application Topology view.	
	Select infrastructure component to display in the Flow Analysis view.	
	Select the application and drill down to Root Cause.	
	Start Packet Capture.	

Table 3.1: Uila Legend

4. Getting Started

This chapter describes the minimum system requirement to install and operate Uila IPM, initial registration steps, and how to install and configure Uila software in vCenter and vSphere environments.

For the following sections, please refer to

- *Uila SaaS Installation Guide*
- *Uila Management Analytics Systems Installation Guide (for On-Premise deployment ONLY)*

for System Requirements, Registration Instructions, and Instructions to install Uila software.

4.1. System Requirements

Always refer to the Uila website for updated system requirements as the first step: https://www.uila.com/products/uila-system-requirements

- Internet Browser for your monitoring console
    - Firefox, Chrome on Windows platform
    - Safari, Firefox, Chrome on OS X platform
    - Firefox, Chrome on CentOS, Ubuntu Linux platform
- VMware version requirements
    - vSphere ESXi 5.5 or higher
    - vCenter Server 5.5 or higher
- VMware® NSX requirement (if Applicable)
    - NSX-V
    - NSX-T™ Data Center
- Uila Virtual Smart Tap (vST) requirements -
    - **vST for On-Premise** -
    - Installed as a guest VM
    - 1 vCPU (1 Core)
    - 1Gb memory
    - 2Gb Storage
    - **vST for Public Cloud –**
    - t2.large for AWS
    - D2s v3 for Azure
- **VIC for VMware, Hyper-v requirements**
    - Installed as a guest VM
    - 1 vCPU (2 Cores)

- 4 GB virtual memory reservation for small deployments of less than 500 VMs,8 GB for 500-1000 VMs, 16GB for more than 1000 VMs
- 8 GB virtual storage, local thin provision

- **VIC for AWS**
  - t2.medium (less than 500 Instances)
  - t2.large (500-1000 Instances)
  - r4.large (1000+ Instances)

- **VIC for Azure**
  - B2S (less than 500 VMs)
  - D2s v3 (500-1000 VMs)
  - A2m v2 (1000+ VMs)

vIC's Virtual Resource allocation (depending on # of VM's monitored) is listed in Table below:

Scope	# of VM Monitored	vCPU	Virtual Memory	Local Storage
Small	**0 ~ 500 VM**	**2 Cores**	**4 GB**	**8 GB**
Medium	**501 ~ 1,000 VM**	**2 Cores**	**8 GB**	**8 GB**
Large	**1,001 ~ 2,000 VM**	**2 Cores**	**16 GB**	**8 GB**

Table 4.1: vIC resource allocation requirements

- Proper vCenter access right is required for vIC to collect structural information and CPU, memory and storage metrics from vCenter, make configuration changes, deploy and setup vST VM. You must have one of the two options pre-configured before vIC deployment:
  1. Full administrative access right (vCenter administrator role), or
  2. Partial administrative access right with the following table of privileges enabled (checked).

Privilege Categories	Privilege Items
Datastore	• Allocate space • Browse datastore • Remove file
Global	• Cancel task
Host	• Local operations->Create virtual machine • Local operations->Delete virtual machine • Configuration → Network Configuration
Network	• Assign network

Resource	• Assign virtual machine to resource pool • Modify resource pool
Scheduled task	• Create tasks • Modify tasks • Remove tasks • Run task
Virtual machine	• Configuration • Guest Operations • Interaction • Inventory • Provisioning • Service configuration • Snapshot management • vSphere replication
dvPort group	• Create • Delete • Modify
vApp	• Add virtual machine • Assign resource pool • Assign vApp • Import

Table 4.2: vCenter access rights table

Visualizing Multiple VMware® vCenter® in a single view

Users can merge **two** separate VMware vCenter and enjoy a single pane of glass into the infrastructure, network and applications. One example of this would be a VDI setup where Virtual desktops are in one vCenter, while the VDI infrastructure servers and backend application servers are hosted in another vCenter. With this new feature, users have the complete end-to-end VDI Application Dependency Mapping visibility across the two vCenters.

- **Network requirements**

- Pre-allocate one IP address for each of the vST's, which can be either static IP address or allocated via DHCP, prior to deployment
- Pre-allocate one static IP address for vIC prior to deployment
- Pre-configure your network to open TCP and UDP ports to allow communications between Uila sub-systems as illustrated in the chart below.
- UMAS –
  - If Cloud UMAS is being used, add ugw1s.uila.com/38.99.127.15 as permitted site on the firewall. Please unblock port 5000 and 443 between vIC and the Uila site.
  - Pre-allocate one static IP if the on premise UMAS is used.
- Make sure port 443 and 902 are open between vIC and Hypervisor hosts

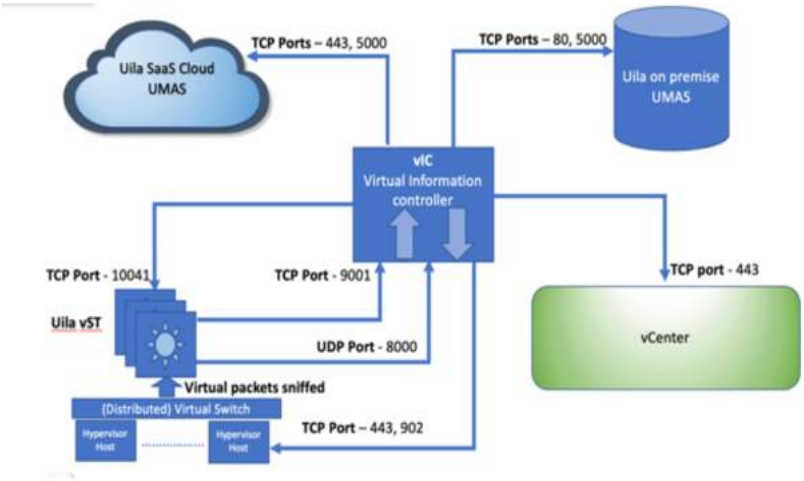

Fig 4.1: Network connection overview for On-Premise Datacenter

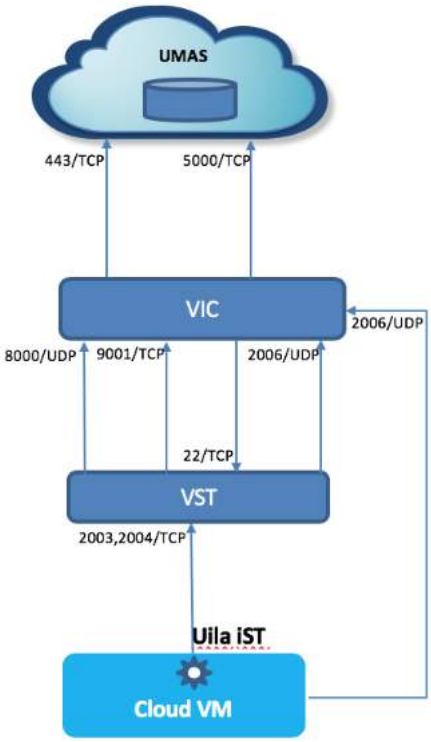

Fig 4.2: Network connection overview for Multi-Cloud Datacenter

5. Baseline

A baseline is a process for monitoring the data center infrastructure's network, compute and storage resources at regular intervals to ensure that the infrastructure which supports business applications are working as intended. It is a process of continually monitoring the key performance indicators to report the health of all applications and its associated data center at a certain point in time. Properly constructing the baseline for your data center, you can obtain the following information:

- Monitor application response time and availability
- Reveal the health state of the infrastructure resources both virtual and physical
- Obtain the current utilization of system resources
- Determine and set alarm thresholds that are unique to your data center operation characteristics
- Alert and identify current system problems that impact Application performance
- Plan for future upgrades and expansions

5.1. Uila Baseline

The baseline methodology is used by Uila extensively. It is the foundation from which *Performance Grades (Infrastructure health performance index) are calculated* and *Alarms* are generated in real time.

Uila maintains a group of *Performance Metrics (See Appendix 15.1)*; for example, Application Response Time, Network Response Time, TCP/IP fatal retry, CPU usage, Memory usage, Disk latency, and many more in its Hadoop data base. Virtual Smart Taps and Virtual Information Manager deployed in user's data center analyze, collect, and transmit these Performance Metrics every minute to Uila Cloud.

Every Metric in per minute interval is compared to a Baseline value for that Metric in real time and a Health Score is calculated based on the formula listed in Table 5.1.

Delta from Baseline	Alarm Severity	Health Score	Color
Less or equal to 5%	**Normal**	75-100	**Green**
Between 5% and 10%, including 10%	**Minor (1)**	50-74	**Yellow**
Between 10% and 20%, including 20%	**Major (2)**	25-49	**Orange**
Above 20%	**Critical (3)**	0-24	**Red**

Table 5.1 – Health score calculations

Uila maintains two kinds of Baseline record for each of Performance Metric monitored;

- **Fixed:** it is a constant value; based on VMware best practices, for example, CPU usage for VM is pre-defined as 80%.
- **Variable:** it is an average of measured metric (per minute) within an hour, i.e. 60 data points. Example of variable metrics are Application Response time, and Network Round Trip time.

During the first day of starting up, current Metrics will be compared to previous hour's value as the default baseline value.

Method of Building Baseline record

Here are the choices you can change how Uila baseline values are defined.

Baseline Metrics	Remarks
Last Hour's value	This is the system default.
Yesterday's value	Select Yesterday's value as the Baseline.
User Configuration option	User selects and locks to a specific week's performance metrics as baseline.

Table 5.2 – Baseline settings

5.2. Health Score and Alarm Definition

Performance Grades are for visual display only and typically color-coded to show the health scores where low score (red) is poor health, and high score (green) is good health. (see Fig 5.1), and are updated every minute.

Here is an example of the Data Center Application Performance summary in color:

Fig 5.1 – Visual display of color-wheel

Alarm is generated based on the performance metric's delta from the baseline. Alarm is generated every 15 minutes by default.

Threshold is defined as the % value that crosses the baseline.

Severity is a user definable indicator to help identify the criticality of the performance metrics monitored to alert user if an entity or entities in his/her data center infrastructure is (are) about to impact the Application's performance.

Delta from Baseline	Alarm Severity	Health Score	Color
Less or equal to 5%	**Normal**	75-100	Green
Between 5% and 10%, including 10%	**Minor (1)**	50-74	Yellow
Between 10% and 20%, including 20%	**Major (2)**	25-49	Orange
Above 20%	**Critical (3)**	0-24	Red

Table 5.3 – Alarm color scheme based on severity

Note: These standard color definitions are applied throughout Uila User Interfaces for consistence and ease of recognition.

6. Managing Your Work from the Console Home Page

Uila console home page is the default infrastructure monitor where the day to day tasks are performed:

- View Application and Infrastructure health dashboard, investigate performance degradation, troubleshooting, and identifying root cause in real time
- Launch additional monitor applications
- Generate reports
- View Syslog
- Change Settings
- Set Preferences
- Go to Full Screen
- See On-line Videos
- Quick Helps

Fig 6.1 – Visual display of dashboard

6.1. Tools Pane

The Tool Pane consists of menu to set up the User profile, and a list of Uila tools for monitoring, report and configuration.

Fig 6.2 – Tools Pane

6.2. Time Matrix Pane

The Time Matrix tool bar allow you to set up a Time Bracket within your timeline horizon where your entire infrastructure performance data are calculated, summarized, compared to prior baseline and displayed in the Monitor pane. You can customize your time window in minutes, hours, or days depending on how you wish to perform real time monitoring, or root cause analysis.

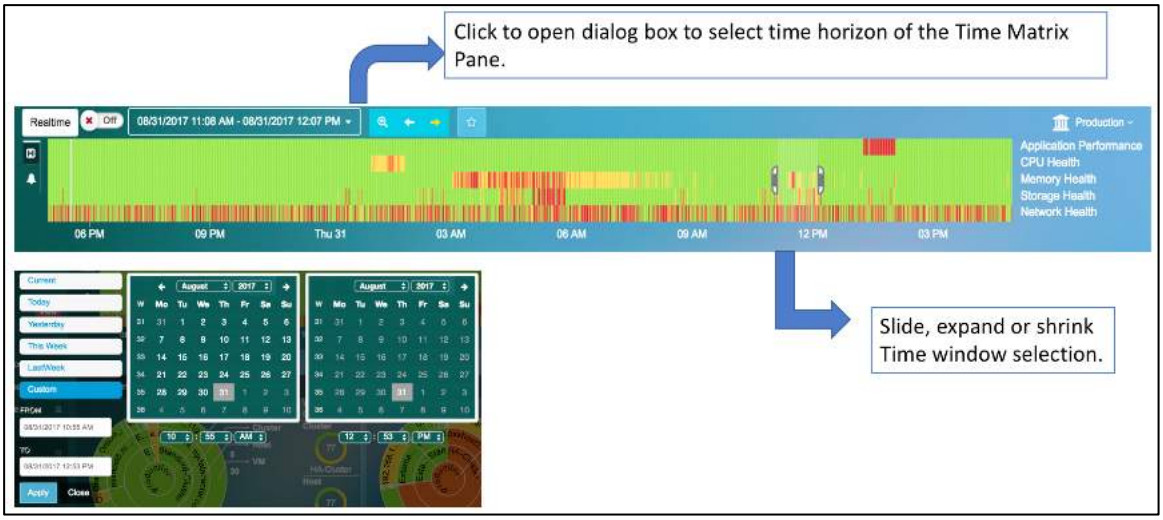

Fig 6.3 – Timeline view

The Time Matrix pane consists of:

- Calendar box to set up time window which you can select between 'Real Time' and 'Time Travel' mode. Select *Current* for Real Time mode.
- Time line window with slide brackets; which can be dragged along the time line to widen or narrow the monitoring window (time range between the brackets)
- Up to five (5) user definable key performance index (KPI) to be monitored. The default KPI are *App Performance*, *CPU Health, Memory Health*, *Storage Health* and *TCP Fatal retry.* The Definition of the first four (4) KPIs are described in the Dashboard chapter in details.

Real Time Mode

In real time mode, all the performance counters are calculated and updated every minute. Typically, you use real time mode to identify root causes of critical applications that exhibit performance degradation in short term, typically in past hours or minutes. System defaults to Real Time mode.

Time Travel Mode

In Time Travel mode, performance data and health measurement metrics are aggregated and calculated based on the Time Bracket you selected. Screen update is stopped. However, data collection continues in real time in the background. Time Travel mode is commonly used for

- Setting infrastructure Baseline to monitor for exceptional events that impact Application performance health. We recommend that you set the larger window bracket what is large enough to obtain a Baseline to represent your infrastructure health that is under normal operation. Common best practice is use a full week that average over several weeks to smooth out exceptional conditions.

- Real time troubleshooting where you may need to travel back in time to look for similar alerting event patterns that impacted performance currently.

6.3. Monitor Pane

The Monitor pane is the working space where Uila tools; such as, Dashboard, Flow Analysis, Application Topology, reports, and other Uila Tool displays its contents as a result of your drill down action. By default, a Dashboard that highlights your infrastructure performance health is displayed after you log in to the system.

6.4. Settings

The settings maintain Uila systems configurations for; (1) vST and vIC software initial installation, and new software updates and upgrades, (2) Interface to physical devices, (3) External systems to receive Alarms.

Here is a list of Configuration Settings Menu:

Menu	Definition
VST Configuration	Use to select which vSwitch(s) in a host to install vST guest VM.
Alarm Configuration	(1) Select Baseline from - Last Hour - Yesterday - Last Week - Any Week since Uila keeps trending records (2) Define Alarm Action. Support delivery alarm by e-mail.
Software Update	List your Uila software version installed, and if new update is available.
vIC Configuration	Contains options to - Monitor external devices - Define custom applications - SNMP configuration for Top of Rack switches - Ignoring certain TCP ports for ART - vIC management (restart, reboot, logging) - Import External Device Address Book Settings
Device Monitoring	Configure Network Device Monitoring capabilities and license usage
Server Configuration	Monitors server settings and license usage
Security Monitoring	Configure threat update intervals and alert filtering
User Experience	Allows user to configure remote sites for end user response time.
Global Configuration	- Define SNMP server IP address, port number, user name, and password to receive the Uila Alarms. - Packet capture configuration

Accounts Management	Allows user to create role-based access control for individual users. AD/LDAP integration can also be enabled to give users access into Uila.
VIC Installation	Step by step instructions to install VIC either the first time, or user wish to deploy VIC in more data centers.

Table 6.1 – Settings meu

6.4.1. User Roles and Privileges

Uila offers three user types –

- Uila Administrator
- Data Center administrator
- Standard User

Here is the comparison of the 3 user roles.

User role	Uila admin	Datacenter admin	Standard User
Number of accounts	Only 1	More than 1	More than 1
Packet Capture	Yes, for all datacenters	Yes, for assigned datacenters	No
vST configuration	Yes, for all datacenters	Yes, for assigned datacenters	No
Alarm Configuration	Yes, for all datacenters	Yes, for assigned datacenters	No
VIC Configuration	Yes, for all datacenters	Yes, for assigned datacenters	No
Device Monitoring	Yes, for all datacenters	Yes, for assigned datacenters	No
Server Monitoring	Yes, for all datacenters	Yes, for assigned datacenters	No
Security Configuration	Yes, for all datacenters	Yes, for assigned datacenters	No
Software Update	Yes, for all datacenters	No	No
User Experience	Yes, for all datacenters	Yes, for assigned datacenters	No
Account Management	Yes, for all datacenters	No	No
Global Configuration	Yes, for all datacenters	Yes, for assigned datacenters	No

Uila administrator can assign pre-built service groups with mission critical servers and applications to a non-administrator user. This would allow a standard user to focus on their relevant multi-tiered applications without having to look at the datacenter as a whole.

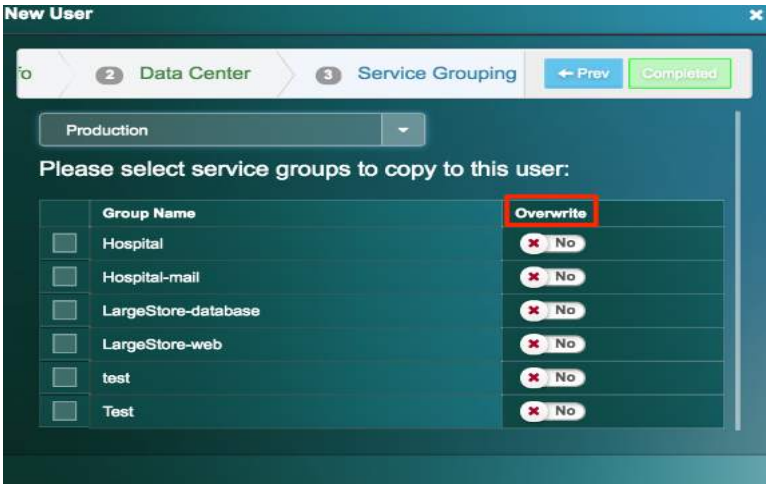

7. Dashboard

Dashboard is the first screen displayed after login. It allows the user to have a unified high-level view of the overall health of the key components in real time and critical alerts that impact the Application performance and Security of the Data Center or Hybrid Cloud deployment.

There are 2 separate Dashboards available: 1) Performance, 2) Security

The Performance Dashboard allows the user to decide on the areas of focus to investigate application slowdown and the issues impacting the Applications performance. The center of the screen shows you the overall health scores in five (5) key areas; ***Application***, ***Network***, ***Storage***, ***CPU*** and ***Memory*** within the infrastructure components, and organized by hierarchical structure relevant to each component in sun burst (color wheel) format.

Fig 7.1a: Performance Dashboard View

The Security Dashboard allows the user to monitor their Cyber Threat status for the entire deployment. This includes getting the overall status for the Cyber Threats that are impacting the Data Center or Cloud deployment, Application Anomalies that have been identified, and finally information on traffic that is exfiltrated (outbound) from the internal VMs.

Fig 7.1a: Security Dashboard View

7.1. Summary of Key Performance Index

The Application and the related infrastructure Health score are monitored according the metric listed in Figure 8.2.

KPI	Metric Monitored	Measurement Method
Application Performance	Application Response Time	Time measured on the server from the arrival of a client request to the transmission of a server response
Network Health	Network Round Trip Time	Packet round trip time spent in the network
	TCP Fatal Retry	TCP re-transmit the same packet for the fourth time or greater
Storage Health	Disk Read Latency	Average amount of time taken process a read command issued from the Guest OS to the virtual machine. The sum of kernelReadLatency and deviceReadLatency in VCDB
	Disk Write Latency	Average amount of time taken processing a Write command issued from the Guest OS to the virtual machine. The sum of kernelWriteLatency and deviceWriteLatency in VCDB
CPU Health	CPU Ready	Percentage of time that the VM was ready, but could not get scheduled to run on the physical CPU due to physical CPU resource congestion
	CPU Usage	Average CPU utilization over all available virtual CPUs in the VM
Memory Health	Swap Wait Time	Time the virtual machine is waiting for memory to be swapped in
	VM Memory Usage	Memory usage as percentage of total configured or available memory

Table 7.1: Infrastructure Health Measurement Metrics and Definitions

7.2. Application Performance Metric

The Application Performance color wheel displays the health of Applications currently running in your data center. The rings present the hierarchical constructs of a virtual Data Center, where you may configure your data center in multiple logical Port Groups. Each Port Group consists of a series of Applications (vApp); such as MySQL, business logics, and web service to perform a specific application function for the end user. These applications depending on the business requirement may run on one or more than VMs.

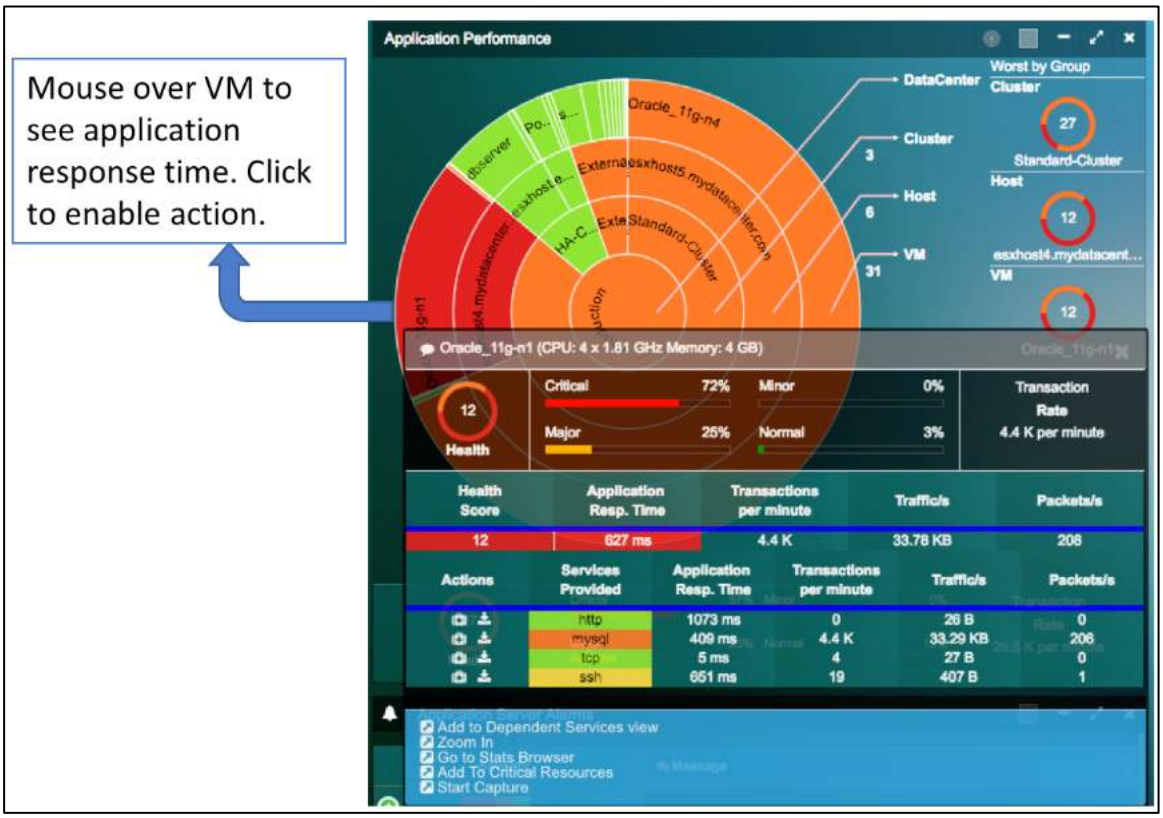

Fig 7.2: Application performance metrics

Application Performance Health Metric

Measurement Metric	Measurement Method	Definition
Application Response Time (in millisecond)	Monitored at packet transaction level	Time measured on the server from the arrival of a client request to the transmission of a server response

Table 7.2: Application performance health metric

Ring Structure and Size Definition

Ring Structure	Color	Size

Ring Center	Data Center	Color represents the averaged Application Performance for the group over the time range selection in the Time Matrix bar.	Application Transaction Volume
Ring 1 (inner ring)	Cluster/Cloud Region		
Ring 2	Host/VPC		
Ring 3 (outer ring)	VM/Instance		

Table 7.3: Ring structure and size definition for Application performance

Full Screen View

To gain a detailed view of the Application Performance Health, click the button, to enlarge the color wheel and add a table view of a complete list of VM, host, cluster monitored, and its associated health score, average application response time, transaction/minute, traffic/second, and packet/second. Each of the column can be sorted by clicking the column header.

Fig 7.3: Application performance detailed view

7.3. Network Performance Metric

The Network Health color wheel displays the health of network with respect to the infrastructure currently running in your data center. The rings present the hierarchical constructs of a virtual Data Center, where it typically structures from TOR Switches, Host, to VM's. Each TOR Switch is connected to a number of Hosts, where one or more VM's resides.

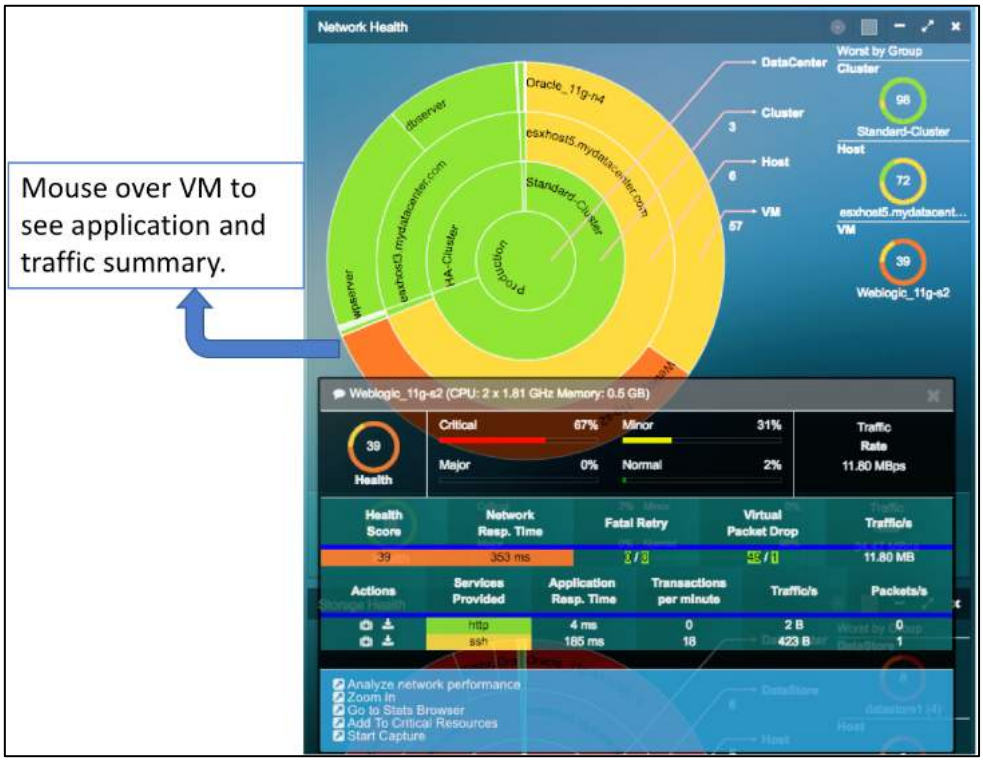

Fig 7.4: Network performance metric

Network Health Metric

Measurement Metric	Measurement Method	Definition
Network Round Trip time (in millisecond)	Monitored at packet level	Packet Round trip time spent in the network
TCP Fatal Retry (in count)	Monitored at packet level	TCP Fatal retry is the TCP packet retransmission for the same packet for the fourth time, which triggers TCP back off algorithm and significant application delay in response.

Table 7.4: Network Health Metric

Ring Structure and Size Definition

Ring Structure		Color	Size
Ring Center	Data Center	Color represents the average weighted Network Health score for each respective group over the time range selection in the Time Matrix bar. See color and baseline	Network Traffic Volume
Ring 1 (inner ring)	Cluster/Cloud Region		
Ring 2	Host/VPC		

Ring 3 (outer ring)	VM/Instance	definition in Time Matrix Bar (Fig 6.3)

Table 7.5: Ring structure and size definition for Network health

7.4. Storage Performance Metric

The Storage Health color wheel displays the health of storage systems currently running in your data center. The rings present the hierarchical constructs of a storage system within your Data Center, where it typically owns multiple Data Stores. Each Data Store groups together a number of Hosts.

Fig 7.5: Storage Health

Storage Health Metric

Measurement Metric	Measurement Method	Definition
Disk Read Latency (in millisecond)	Sourced from vCenter (VCDB)	Time taken to complete a Read command issued from the Guest OS. This Disk Read Latency includes VM kernel Read Latency and Device Read Latency.
Disk Write Latency (in millisecond)	Sourced from vCenter (VCDB)	Same as the above for Write command.

Table 7.6: Storage Health Metric

Ring Structure and Size Definition

Ring Structure		Color	Size
Ring Center	Data Center	Color represents the average weighted Storage Health score for each respective group over the time range selection in the Time Matrix bar. See color and baseline definition in Time Matrix Bar (Fig 6.3)	Number of Storage I/O Operations
Ring 1 (inner ring)	Data Store		
Ring 2	Host/VPC		
Ring 3 (outer ring)	Virtual Disk		

Table 7.7: Ring structure and size definition for Storage Health

Full Screen View

To gain a complete detail view of the Storage Performance Health, click the  button, to enlarge the color wheel and add a table view of a complete list of VM, host, data store monitored, and its associated health score, read latency, read IOPS, write latency, write IOPS. Each of the column can be sorted by clicking the column header.

Fig 7.8: Storage performance full screen view

Storage Disk usage charts and alerts: Users now have access to new circle packing views and tables to visualize storage disk usage and capacity.

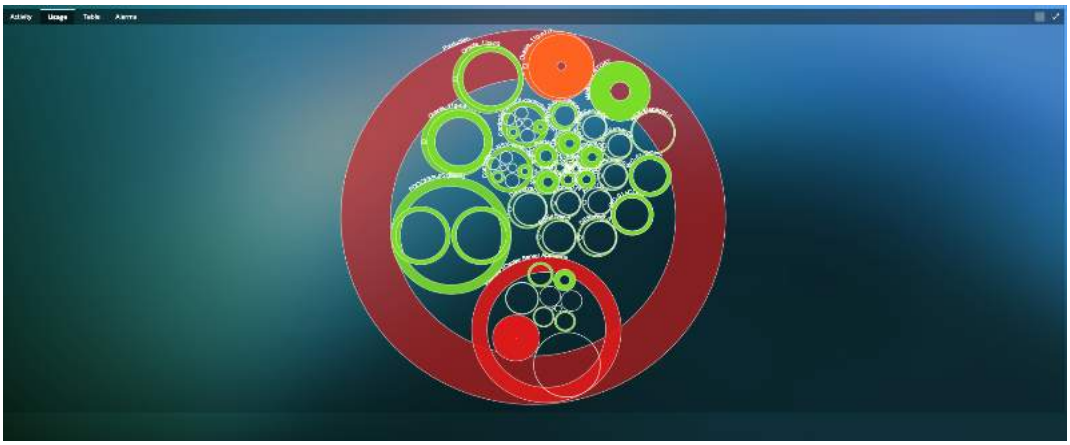

7.5. CPU Performance Metric

The CPU Health color wheel displays the performance of all CPU in your Hosts with respect to the infrastructure currently running in your data center. The rings present the hierarchical constructs of a virtual Data Center, where it typically structures to form a number of Cluster of Hosts, under which one or more VM's resides.

Fig 7.9: CPU metric

CPU Health Metric

Measurement Metric	Measurement Method	Definition
CPU-Ready (%)	Sourced from vCenter (VCDB)	Percentage of time that the VM was ready to run, but could not get scheduled to run on the physical CPU due to physical CPU resource congestion.
CPU Usage (%)	Sourced from vCenter (VCDB)	CPU usage is the percentage of active CPU to total configured CPU.

Table 7.8: CPU Health Metric

Host CPU Metric Calculation

Measurement Metric	Normal (Green)	Minor (Yellow)	Major (Orange)	Critical (Red)
CPU-Ready (%) (X = CPU.Ready/ # of pCPU)	X < 6,000 ms (10% per 1 min)	6,000 ms <= X < 9,000ms (10% ~ 15%)	9,000 ms <= X < 15,000ms (15% ~ 25%)	X >= 15,000 ms (>= 25%)
Y=CPU Usage (%)	Y <= 80%	80% < Y <= 85%	85% < Y <= 90%	Y > 90%

Table 7.9: Host CPU Health Metric Calculations

Note:

Host CPU Ready Time = Sum of all pCPU's Ready Time.

VM CPU Metric Calculation

Measurement Metric	Normal (Green)	Minor (Yellow)	Major (Orange)	Critical (Red)
CPU-Ready (%) (X = CPU.Ready/ # of vCPU)	X < 3,000 ms (5% per 1 min)	3,000 ms <= X < 6,000ms (5% ~ 10%)	6,000 ms <= X < 12,000ms (10% ~ 20%)	X >= 12,000 ms (>= 20%)
Y=CPU Usage (%)	Y <= 80%	80% < Y <= 85%	85% < Y <= 90%	Y > 90%

Table 7.10: VM CPU Health Metric Calculations

Ring Structure and Size Definition

Ring Structure		Color	Size
Ring Center	Data Center	Color represents the average weighted CPU Health score for each respective group over the time range selection in the Time Matrix bar. See color and baseline definition in Time Matrix Bar (Fig 6.3)	
Ring 1 (inner ring)	Cluster/Cloud Region		Physical CPU capacity (MHz)
Ring 2	Host/VPC		Physical CPU capacity (MHz)
Ring 3 (outer ring)	VM/Instance		

Table 7.11: Ring structure and size definition for CPU Health

Full Screen View

To gain a complete detail view of the Storage Performance Health, click the  button, to enlarge the color wheel and add a table view of a complete list of VM, host, data store monitored, and its associated Health score, Application Response Time, Usage %, Usage MHz, CPU Ready. Each of the column can be sorted by clicking the column header.

Fig 7.10: CPU performance full screen view

7.6. Memory Performance Metric

The Memory Health color wheel displays the performance of all memory arrays in your Hosts with respect to the infrastructure currently running in your data center. The rings present the hierarchical constructs of a virtual Data Center, where it typically structures to for a number of Cluster of Hosts, under which one or more VM's resides.

Fig 7.11: Memory performance metric

Memory Health Metric

Measurement Metric	Measurement Method	Definition
Swap Wait time (milliseconds)	Sourced from vCenter (VCDB)	Time the virtual machine is waiting for memory pages to be swapped in.
Memory Usage (%)	Sourced from vCenter (VCDB)	VM Memory usage is the percentage of active memory to total configured memory. Host and Cluster Memory Usage is the percentage of consumed memory (including VMkernel and Guest VMs) to physical memory capacity.
Swap-in Rate (kbps)	Sourced from vCenter (VCDB)	Average amount of memory (kbps) swapped in from disk into memory for VM to run.

Table 7.12: Memory Health metric

Host Memory Metric Calculation

Measurement Metric	Normal (Green)	Minor (Yellow)	Major (Orange)	Critical (Red)
Swap-Wait (%) (X = Swap-Wait/ # of pCPU)	X < 6,000 ms (10% per 1 min)	6,000 ms <= X < 9,000ms (10% ~ 15%)	9,000 ms <= X < 15,000ms (15% ~ 25%)	X >= 15,000 ms (>= 25%)

Table 7.13: Host Memory Health calculations

Where:

X=CPU.SwapWait /# pCPU (ref %SWPWT in ESXTOP)

VM CPU Metric Calculation

Measurement Metric	Normal (Green)	Minor (Yellow)	Major (Orange)	Critical (Red)
Swap-Wait (%) (X = Swap-Wait/ # of vCPU)	X < 3,000 ms (5% per 1 min)	3,000 ms <= X < 6,000ms (5% ~ 10%)	6,000 ms <= X < 12,000ms (10% ~ 20%)	X >= 12,000 ms (>= 20%)
Y= Mem Usage (%)	Y <= 70%	70% < Y <= 75%	75% < Y <= 85%	Y > 85%

Table 7.14: VM Memory Health calculations

Note:

VM CPU Swap Wait Time = Sum of all vCPU's Swap Wait Time.

VM Mem Usage = Active / Virtual machine configured size.

Ring Structure and Size Definition

Ring Structure		Color	Size
Ring Center	Data Center	Color represents the average weighted MEMORY Health score for each respective group over the time range selection in the Time Matrix bar. See color and baseline definition in Time Matrix Bar (Fig 6.3)	
Ring 1 (inner ring)	Cluster/Cloud Region		Physical MEMORY capacity (MHz)
Ring 2	Host/VPC		Physical MEMORY capacity (MHz)
Ring 3 (outer ring)	VM/Instance		

Table 7.15: Ring structure and size definition for Memory Health

The consolidation ratio is a measure of the number of VMs placed on a physical machine. ESX Server's over commitment technology is an enabling technology allowing users to achieve a higher consolidation ratio, thus reducing the total cost of operation. Over commitment is the ability to allocate more virtual resources than available physical resources. ESX Server offers users the ability to overcommit memory and CPU resources on a physical machine.

Full Screen View

To gain a complete detail view of the Storage Performance Health, click the  button, to enlarge the color wheel and add a table view of a complete list of VM, host, data store monitored, and its associated Health score, Application Response Time, Usage %, Active, CPU Swap Wait. Each of the column can be sorted by clicking the column header.

Fig 7.12: Memory performance full screen view

8. Application

8.1. Dependency Mapping

Application Analysis provides you a visual view of all virtual Application (vAPP) service chains within your data center in real time. Applications within a defined Port Group are grouped together to help you quickly identify how each Application and its associated VM is communication with each other. It shows the health of each VM by calculating the average application response time of the VM server.

The application dependency map will also extend beyond

Application Analysis view is directly launched from the Tool Pane menu, and it consists three tabs (views):

- Topology Map view: See complete view of all application servers inside a vCenter
- Dependent Services view: See application service chaining. Multiple views can be customized
- Table view: Organize in table view to sort by performance grade of the VM. Refer to Chapter 7.2 Application Performance Metrics for details.

8.1.1.Topology Map View

You can use Topology Map view to see all application servers (VM's) organized by Port Group (VLAN) view in a glance, and how they communicate with each other. This view is particularly useful for

- Revealing how and if Port Groups (VLAN) are interconnected
- Showing each application service performance by its response time and transaction load on the associated VM's
- Identifying any orphan VM's (VM's are standalone without communication with any other VM), which are the result of misconfiguration.
- Identifying any application services performance degradation and pinpoint the root cause quickly.

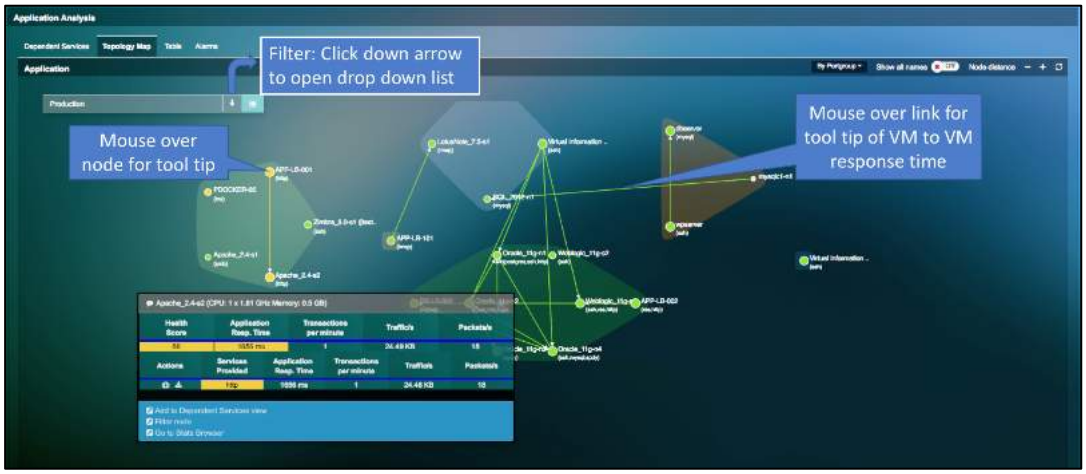

Fig 8.1: Application topology map view

Symbol	Definition	Mouse Over Information	Click Action
	Application VM name with list of protocol identified.	Highlight connections between this Application VM and neighbor VM's	Select one of the protocols to identify the root cause of slow response time

		Show a list of active Application protocols and associated response time	
→	Traffic flow between Applications	Displays average transaction response time between two VM's for each of the application service running.	None
	Find Root Cause for Application issue	None	Click to Root Cause view
	Packet Capture Network Traffic for the selected Application	None	Click to start packet capture

Table 8.1: Symbols, definition, information and action

You can visualize the properties of the VM/server, from the properties menu option, when you click on any VM/server.

8.1.2. Dependent Service View

Dependent service view is particularly useful when you have a large number of application servers (VM's) that are crowding your screen, and you are interested in only those critical application service chaining that runs your mission critical business applications. There is no practical limit of how many Dependent Service view you can create and customized.

To create a Dependent Service view, follow these steps:

1. Find VM that is the beginning of your critical service chaining, click to show the VM health summary
2. Select and click the "Add to Dependent Services View"

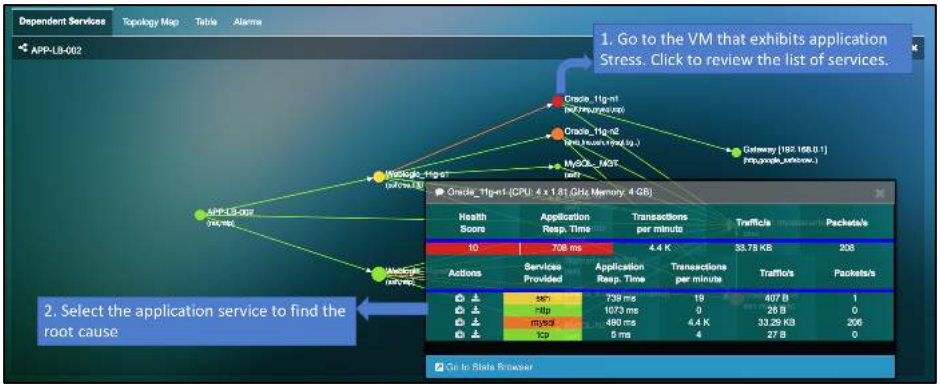

Fig 8.2: Application topology to dependent service view

A new Dependent Service is created, see example below, and notes the steps of finding the root cause of application performance degradation.

8.1.3. Service Filter

The function in application dependency mapping filters the Dependency Mapping window to display only the selected service or application. This allows the user to focus on the services or applications that needs to be monitored or troubleshoot for user complaints.

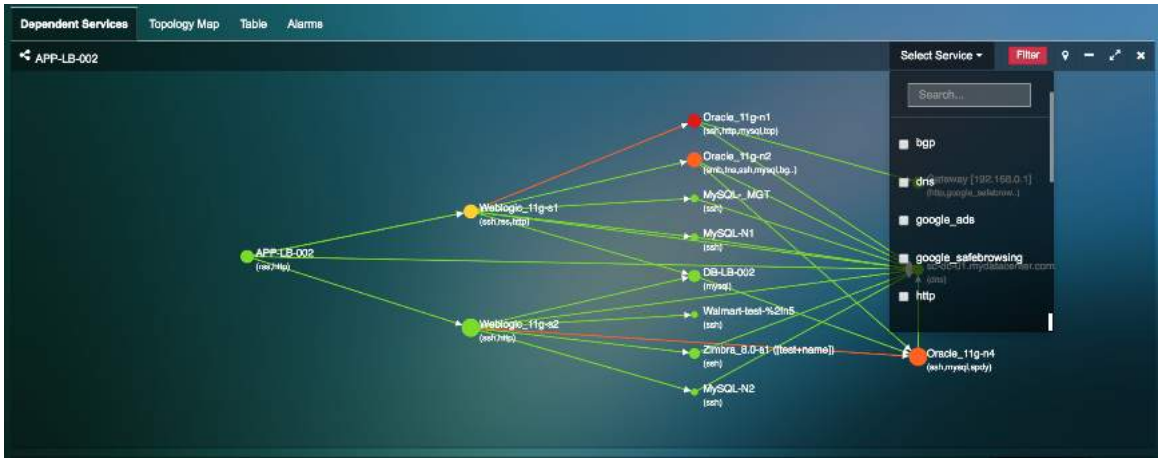

Fig 8.3: Service Filter in Application Dependency Mapping

8.1.4. Multi-Cloud Application Dependency Mapping

Uila's Multi-Cloud Application Dependency Maps provides the user with the ability to see the application dependencies across the cloud boundaries. Uila makes it easy to visualize application on the cloud and their dependencies to on-premise servers.

8.1.5. Resolve Gateway

The “Resolve Gateway” button removes the gateway from showing up on the Application Dependency map. This can be helpful when the user wants to see the direct dependencies of servers within the environment.

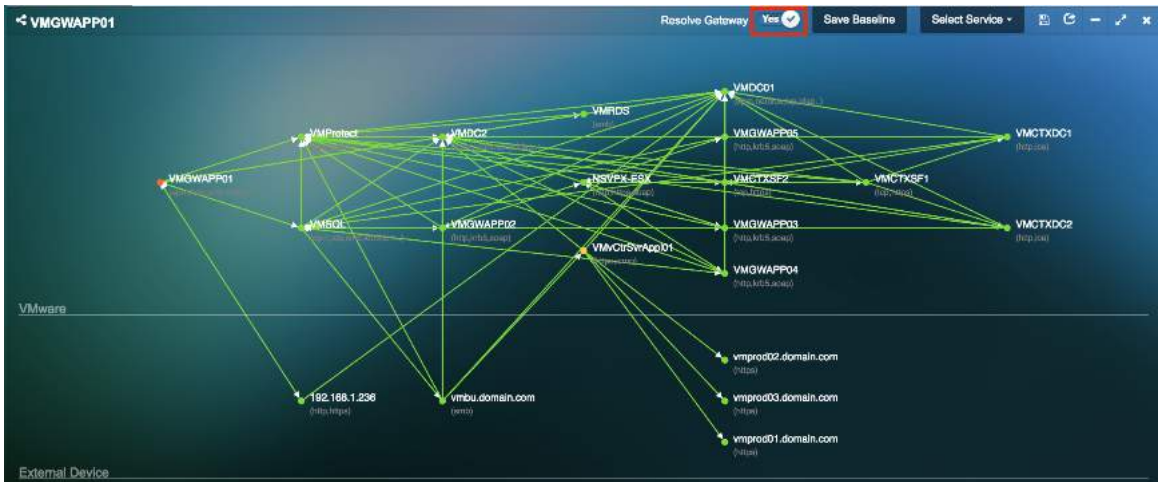

8.1.6. Change control Monitoring and Baselining

Uila’s change control monitoring and baseline feature provides the user with the ability to baseline the application dependency map during the normal course of operation. The application can be baselined and compared to the application dependencies to any given time period. With the change monitoring capability, users can stay on top of all changes in the applications, servers delivering those applications and the interdependencies in the environment, including new entrants and exits.

1) The baseline can be set by selecting the appropriate range of normal functioning and clicking the “Save Baseline” button.

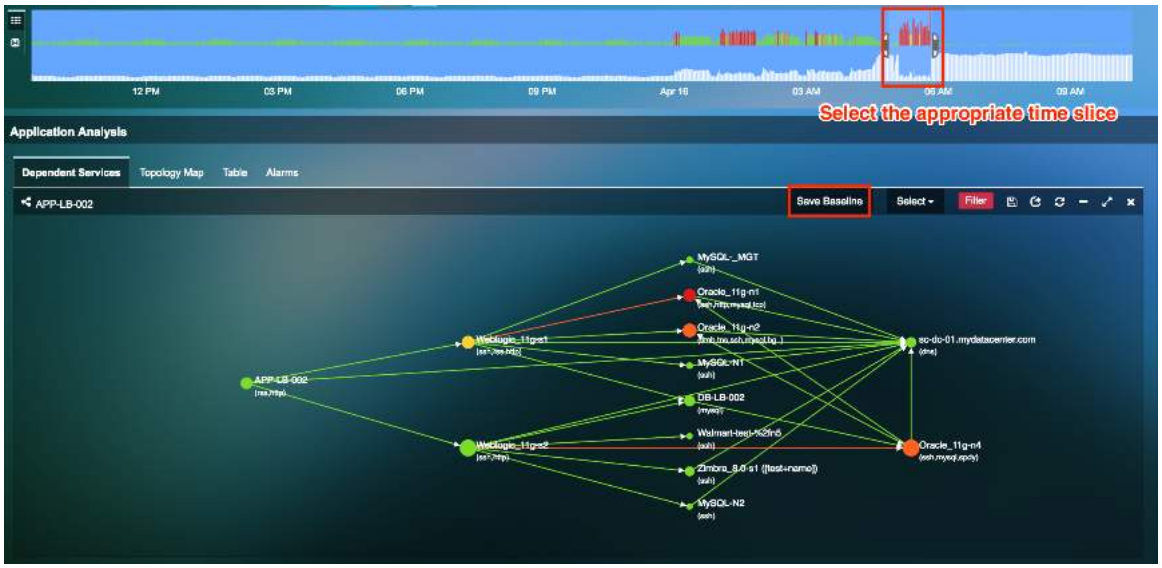

2) When the user moves the time slider to another timeframe, Uila will now report all the changes in connection.

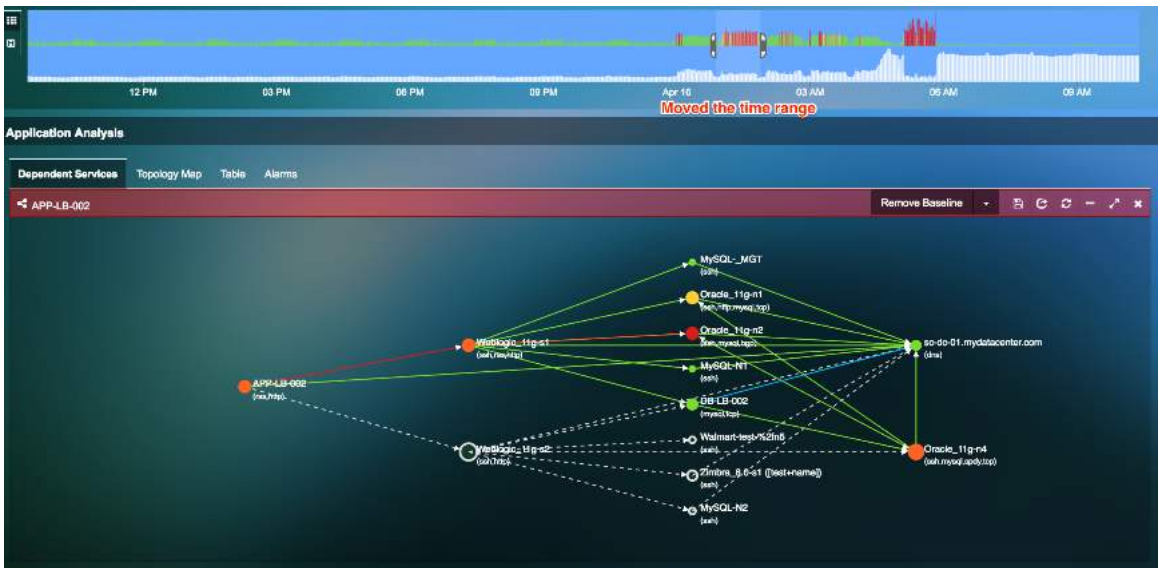

Dotted gray line - The dotted gray line seen on the map indicates all the missing inter-connections in comparison to the baseline.

Blue Line – The solid blue line indicates any new dependencies and inter-connections between the individual VM's.

8.1.7. Display External IP addresses and MAC addresses on the Application

External devices may include physical servers, VMs in a separate Data Center, gateways, firewall, load balancer, client devices, VM running in any cloud provider's platform, network switches, etc. Now the user can display those external devices in their Application Dependency Map by entering its IP address. This is enabled from the Settings→VIC configuration menu.

- **Manually display External Device by IP**

1) Go to Settings → VIC configuration
2) In order to add a new External Device by IP, click New

3) Add the fields –

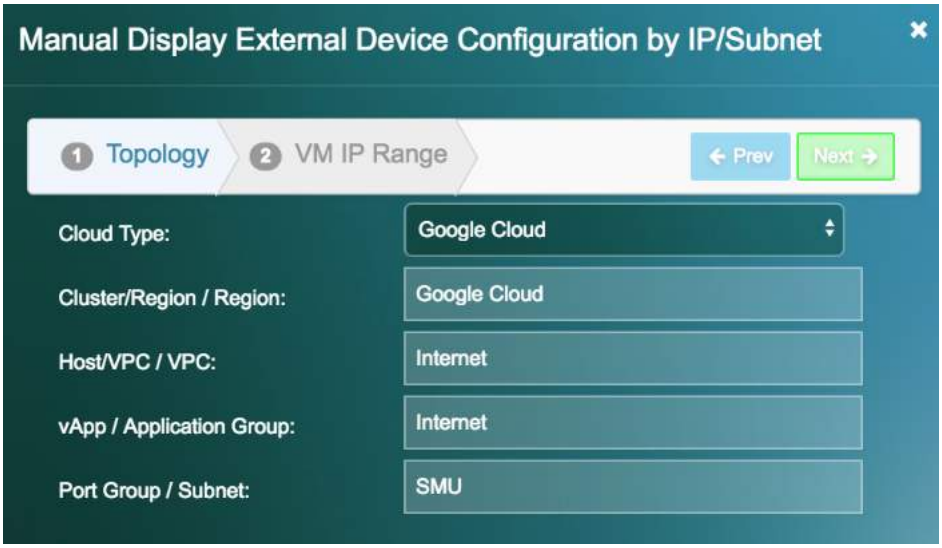

4) Select the IP ranges –

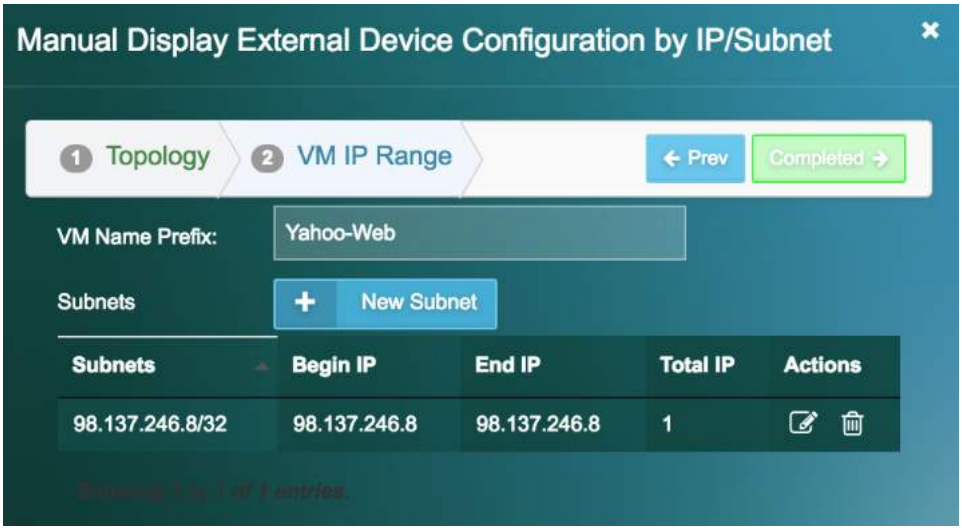

5) Now you will see these devices appear on the Application Dependency Map

- **Manually Display External Device by MAC**

1) Go to Settings → VIC configuration
2) In order to add a new External Device by MAC, click New
3) Click on “New MAC Address” to add the device –

4) Add the MAC’s

5) Now you will see the device appear on the Application Dependency Map

8.1.8.Application dependency map and server topology map export

Users can export the application dependency map and server topology map into an excel spreadsheet. A common use case for this export is it can be used for datacenter pre-migration assessments to the Hybrid Cloud.

1) To export the application dependency map, click on the “Export Application map”

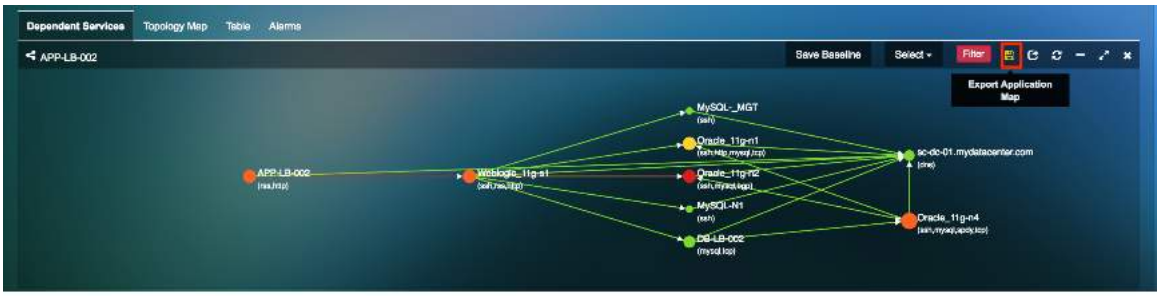

2) The CSV export provides the user with the excel sheet to help identify the various inter-dependencies and the capacity of individual virtual machines. There are 2 sections in the excel sheet; Dependency and Capacity.
    a. Dependency – This provides us with all the inter-connections between different servers; the source, through the gateway and the destination. It also provides us with the port numbers and the applications.
    b. Capacity – This provides information on each server, the number of CPU cores and the memory allocated to each server.

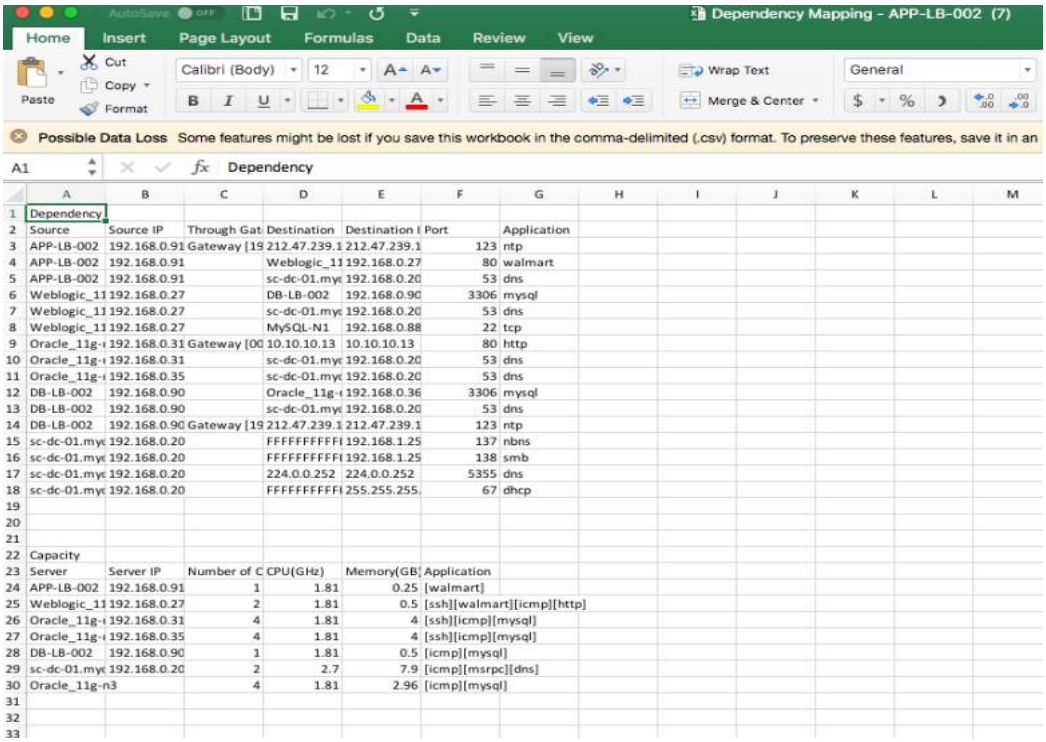

	A	B	C	D	E	F	G
1	Dependency						
2	Source	Source IP	Through Gat	Destination	Destination I	Port	Application
3	APP-LB-002	192.168.0.91	Gateway [19	212.47.239.1	212.47.239.1	123	ntp
4	APP-LB-002	192.168.0.91		Weblogic_11	192.168.0.27	80	walmart
5	APP-LB-002	192.168.0.91		sc-dc-01.my	192.168.0.20	53	dns
6	Weblogic_11	192.168.0.27		DB-LB-002	192.168.0.90	3306	mysql
7	Weblogic_11	192.168.0.27		sc-dc-01.my	192.168.0.20	53	dns
8	Weblogic_11	192.168.0.27		MySQL-N1	192.168.0.88	22	tcp
9	Oracle_11g-r	192.168.0.31	Gateway [00	10.10.10.13	10.10.10.13	80	http
10	Oracle_11g-r	192.168.0.31		sc-dc-01.my	192.168.0.20	53	dns
11	Oracle_11g-r	192.168.0.35		sc-dc-01.my	192.168.0.20	53	dns
12	DB-LB-002	192.168.0.90		Oracle_11g-r	192.168.0.36	3306	mysql
13	DB-LB-002	192.168.0.90		sc-dc-01.my	192.168.0.20	53	dns
14	DB-LB-002	192.168.0.90	Gateway [19	212.47.239.1	212.47.239.1	123	ntp
15	sc-dc-01.my	192.168.0.20		FFFFFFFFFF	192.168.1.25	137	nbns
16	sc-dc-01.my	192.168.0.20		FFFFFFFFFF	192.168.1.25	138	smb
17	sc-dc-01.my	192.168.0.20		224.0.0.252	224.0.0.252	5355	dns
18	sc-dc-01.my	192.168.0.20		FFFFFFFFFF	255.255.255.	67	dhcp
19							
20							
21							
22	Capacity						
23	Server	Server IP	Number of C	CPU(GHz)	Memory(GB)	Application	
24	APP-LB-002	192.168.0.91	1	1.81	0.25	[walmart]	
25	Weblogic_11	192.168.0.27	2	1.81	0.5	[ssh][walmart][icmp][http]	
26	Oracle_11g-r	192.168.0.31	4	1.81	4	[ssh][icmp][mysql]	
27	Oracle_11g-r	192.168.0.35	4	1.81	4	[ssh][icmp][mysql]	
28	DB-LB-002	192.168.0.90	1	1.81	0.5	[icmp][mysql]	
29	sc-dc-01.my	192.168.0.20	2	2.7	7.9	[icmp][msrpc][dns]	
30	Oracle_11g-n3		4	1.81	2.96	[icmp][mysql]	
31							
32							
33							

8.1.9. Automated Application dependency map generation for VDI

For VMware Horizon® versions 6 or higher, Uila automatically generates the Application Dependency Map which can display the different tiers of the entire VDI environment, including thin clients, VDI desktops, as well as critical infrastructure components such as the Connection server, Domain Controller, etc. With this automatically generated map, Uila users are able to automatically highlight the bottlenecks in their VDI environment.

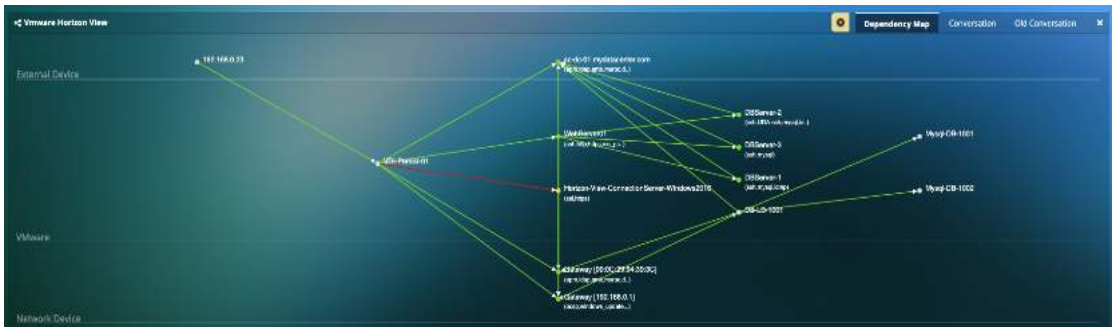

8.1.10. Conversation Map

Users can visualize the applications or services in use on the VMs. For example, this can be very helpful to visualize applications in use on the VDI desktops.

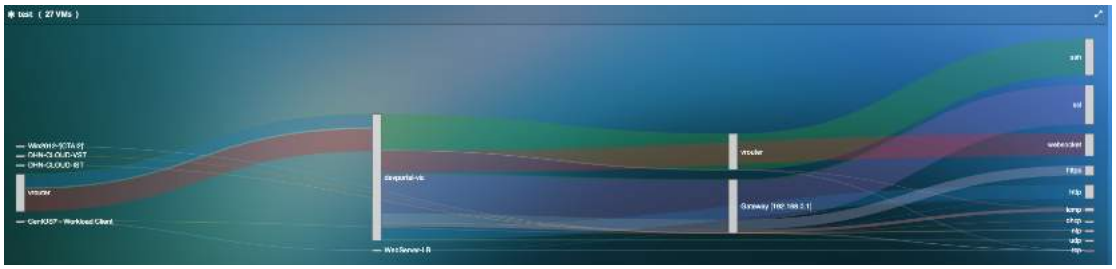

8.2. Transaction Analysis

Transaction Analysis provides deep insights and analytics into web and database application (HTTP, MySQL, Oracle and PostGreSQL) performance. This is done by collecting application response times through the network and by reading transaction codes and queries from the packet. The goal is to provide deeper insights into client and server errors so that the issues can be narrowed down and mitigated.
Transaction Analysis does not require any additional configurations. The vST can immediately identify the type of application traffic and its status codes and query's by parsing through its header file.

This feature provides the users with an overview and individual server view. The overview provides a quick summary of all status codes and queries seen within the entire datacenter. The server view provides a summary of status codes and queries seen by individual servers.

8.2.1.Overview page

Choose the Database you would like to view statistics on using the dropdown –

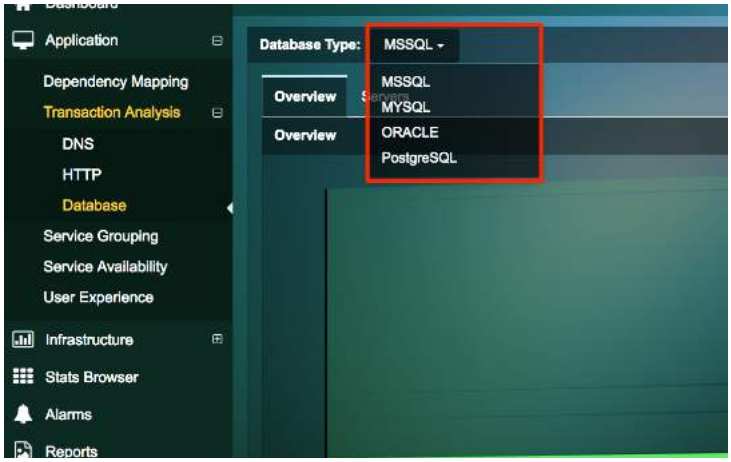

Overview page has 3 components –

- **Ribbon View -**

This view provides the user with a visual representation of the different queries and statuses of individual servers.

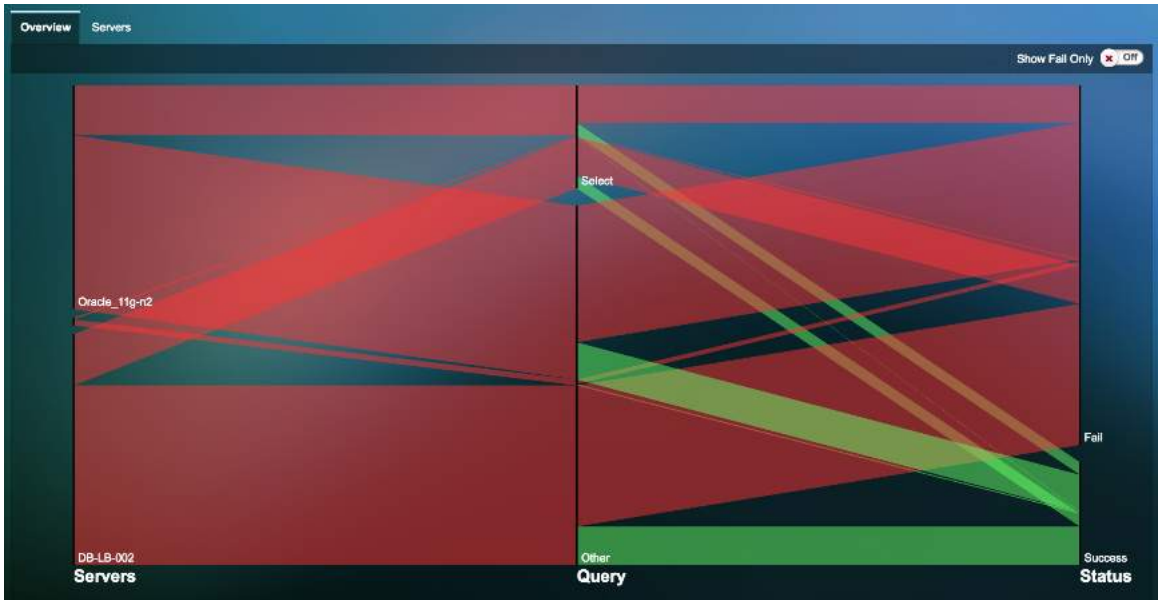

Figure 8.4: Ribbon View

The user can hover over the ribbon to view the server's transaction volume based on queries and the status codes.

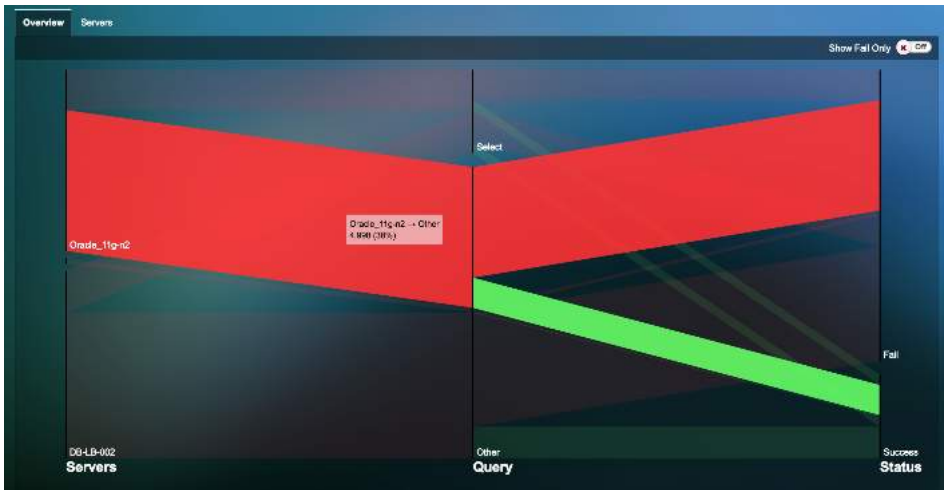

Figure 8.5: Hover over to view details

- **Status code statistics -**

The status code statistics displays the number of status code responses collected. Each vertical bar on the graph represents the number of responses collected per minute.

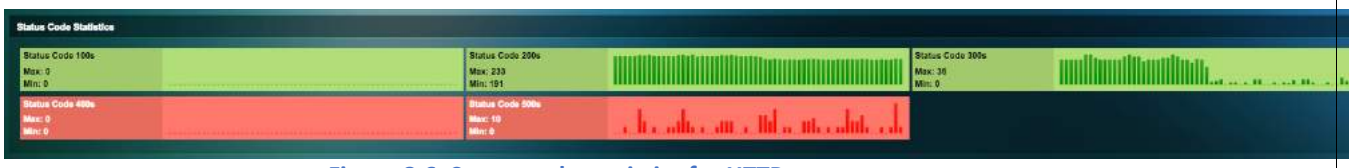

Figure 8.6: Status code statistics for HTTP

Figure 8.7: Status code statistics for MySQL and Oracle

Status Code #	Function
100's	Informational response –continue, switch protocols, processing
200's	Success response – OK, created, Accepted
300's	Redirection response – found, moved permanently, use proxy
400's	Client errors – bad request, forbidden, not found
500's	Server errors – bad gateway, gateway timeout, service unavailable

Table 8.2: Status codes and their function for HTTP

- **Query statistics -**

Displays the application response times and counts per minute for various HTTP (GET, POST, HEAD) and SQL (INSERT, UPDATE, DELETE) queries.

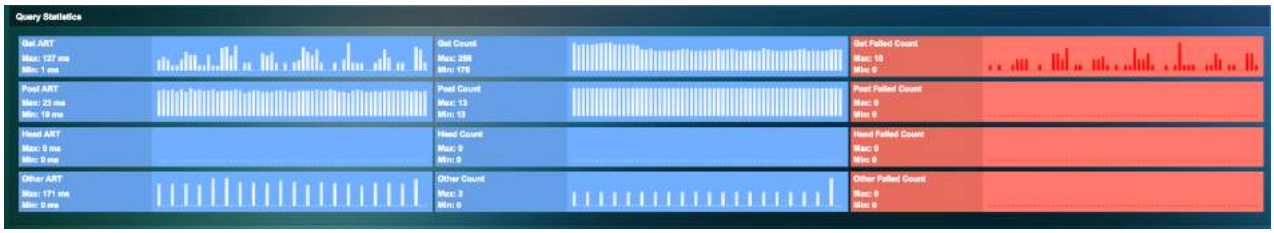

Figure 8.8: Query statistics for HTTP

Query	Function
GET	Gets information from the webserver
POST	Sends data to a webserver
HEAD	Checks if a webserver exists

Table 8.3: Query statistics for HTTP

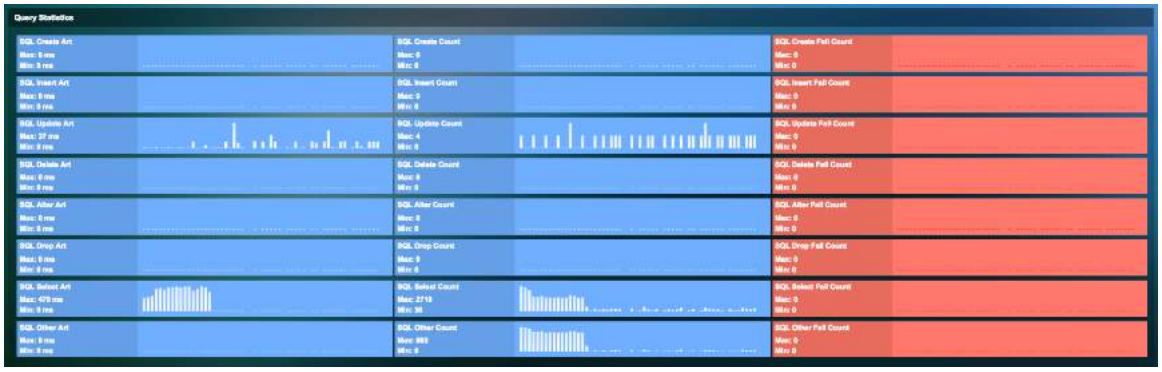

Figure 8.9: Query statistics for MySQL and Oracle

Query	Function
CREATE	Creates a table
INSERT	Inserts into table
UPDATE	Modifies existing records in a table
DELETE	Deletes existing records within table
ALTER	Adds, deletes or modifies columns in existing table
DROP	Drops an exisiting table from a schema
SELECT	Select a database where operations are performed

Table 8.4: Query statistics for MySQL and Oracle

- **Network statistics -**

Displays network specific information such as transaction volume, network delay time and retry rates for the HTTP or database applications.

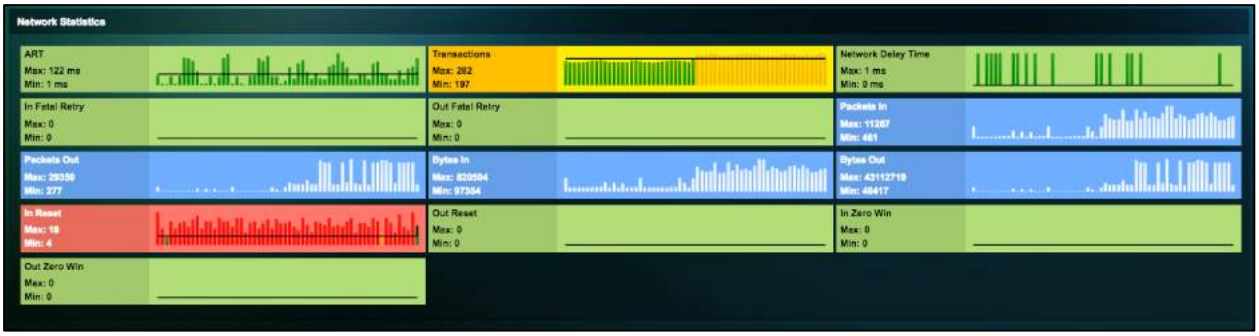

Figure 8.10: Network statistics

Statistic	Function
ART (ms)	Provides the application response times per minute
Transaction	Number of transactions per minute
Network Delay TIme	Network delays per minute
In Fatal Retries	Number of fatal retries inbound per minute
Out Fatal Retries	Number of fatal retries outbound per minute
Packets In	Packets inbound per minute
Packets Out	Packets outbound per minute
Bytes In	Bytes inbound per minute
Bytes Out	Bytes outbound per minute

Table 8.5: Network Statistics

8.2.2. Server page

The server page provides an insight into the individual servers providing the service. Each server's queries and statues are displayed individually to help understand the problematic services.

Figure 8.11: Server page view

Users can further drill down and get more information on their status, queries, network, usage, dependent services and process monitoring.

Figure 8.12: Drill down into server

8.3.3. Transaction Logging

In order to view transaction analysis, the user must redeploy the VST. Once the VST is redeployed -

1) On the Uila dashboard settings -> VST configuration
2) Click on configuration for the VST you would like to enable transaction logs
3) Check “Enable Transaction Analysis” box.

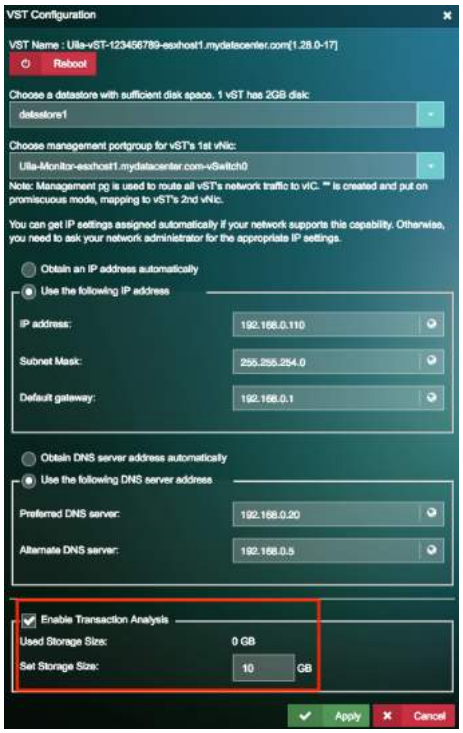

Figure 8.13: Enable Transaction analysis

Once transaction analysis is enabled, you can view the transaction logs on your Transaction analysis view.

You can click on any of the bold underlined hyperlink to view more information on the individual transactions.

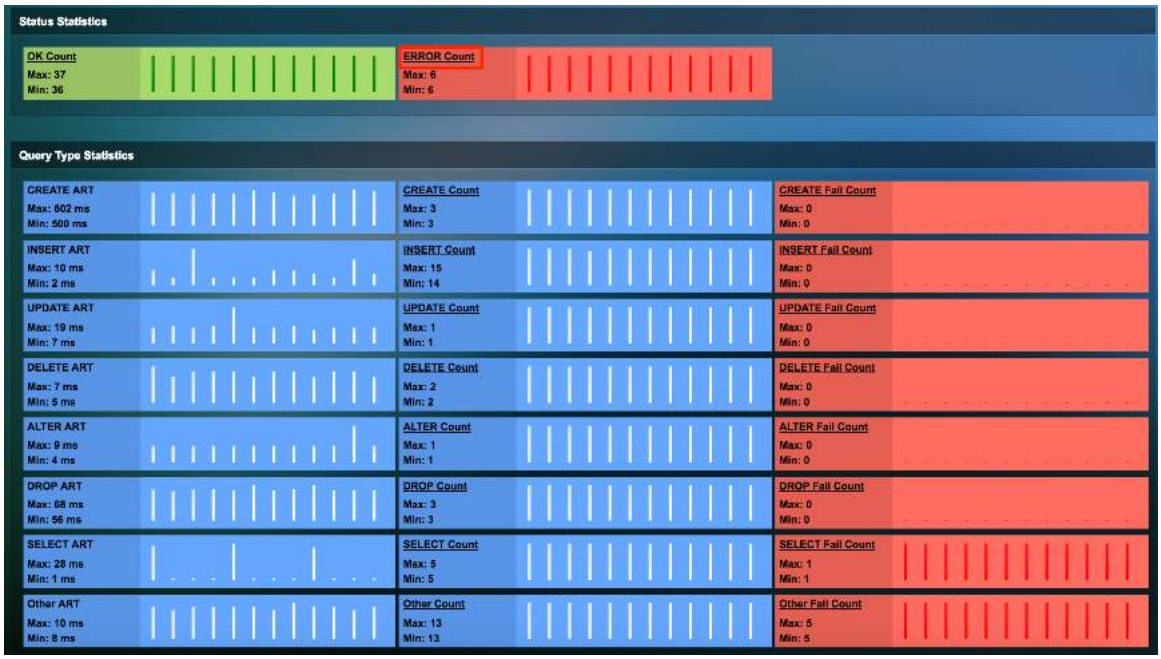

Figure 8.14: Click on the underlined text to view transaction analysis

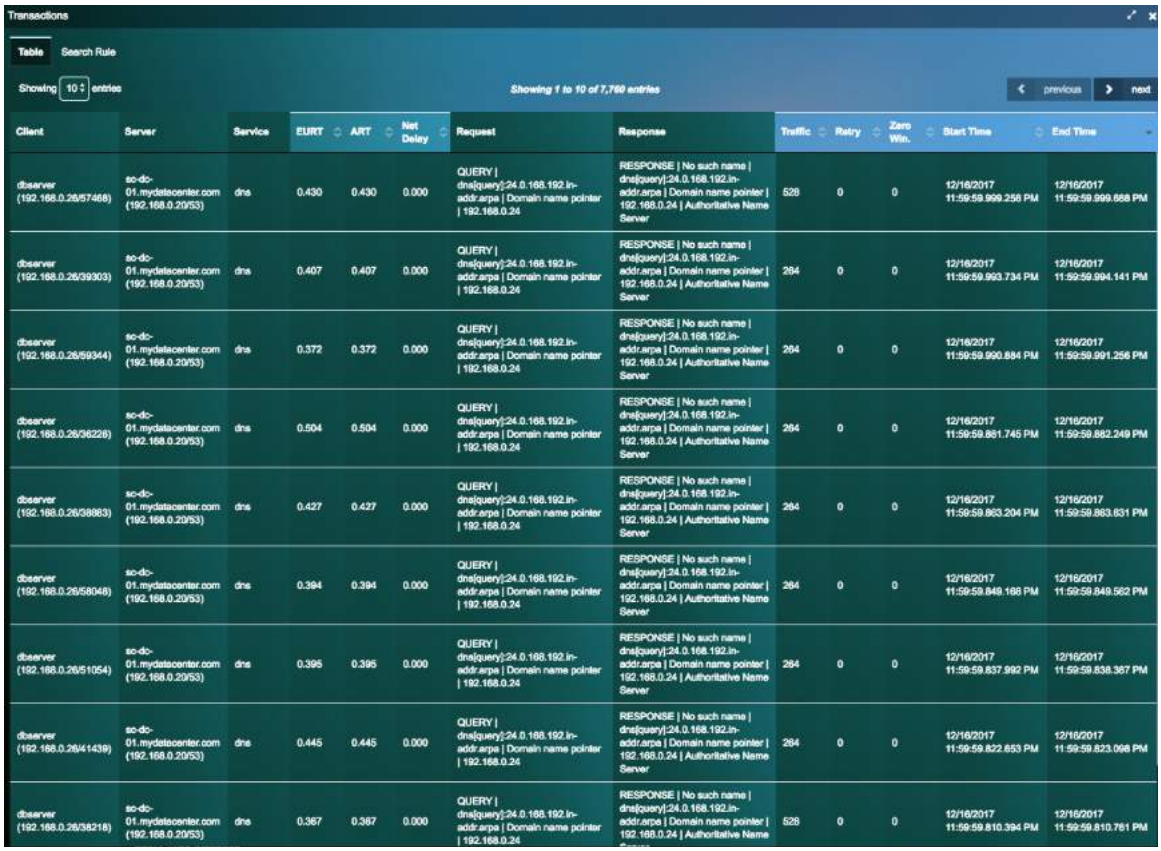

Figure 8.15: Transaction Logs

- **Transaction search analysis -** Users can now search for specific metadata (text) across a multi-tier application chain. For example, you can search for any specific keyword across the datacenter transactions.

  The user can search for specific transactions using the search view –

Figure 8.16: Search function for transactions

Within the search functionality, the “green +” represents AND and “blue +” represents OR.

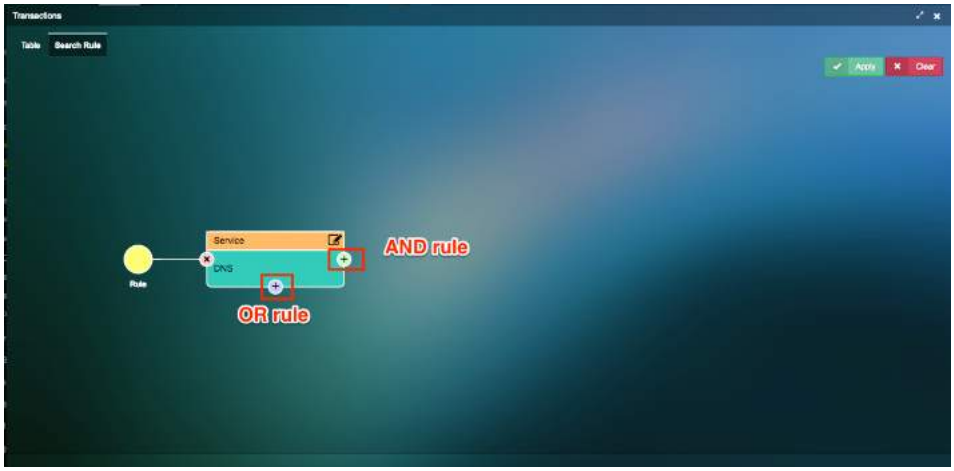

Figure 8.17: Search function

The rules can be setup based on 22 criteria’s as shown in the picture below.

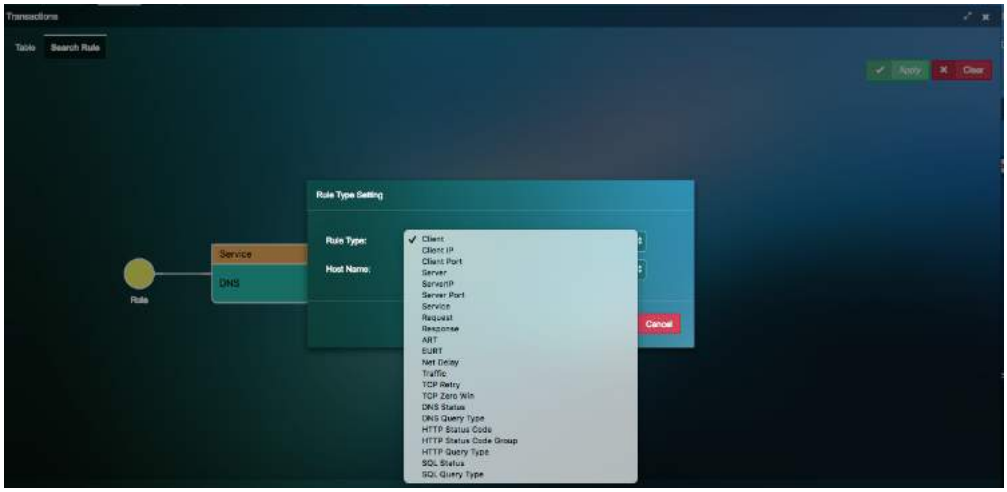

Figure 8.18: Search Criteria

You can also configure # of transaction records exported in CSV for Transaction Analysis.

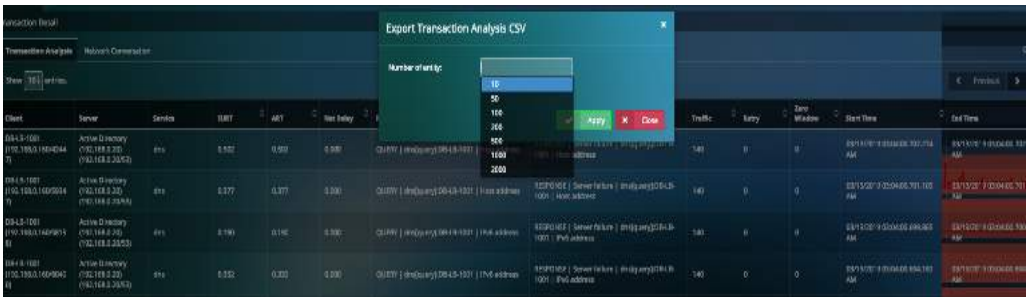

- **Network Conversation**
  Network conversation view provides a list of Network conversations between clients and servers along with their End-User Response time, Network Response time and Application Response Time.

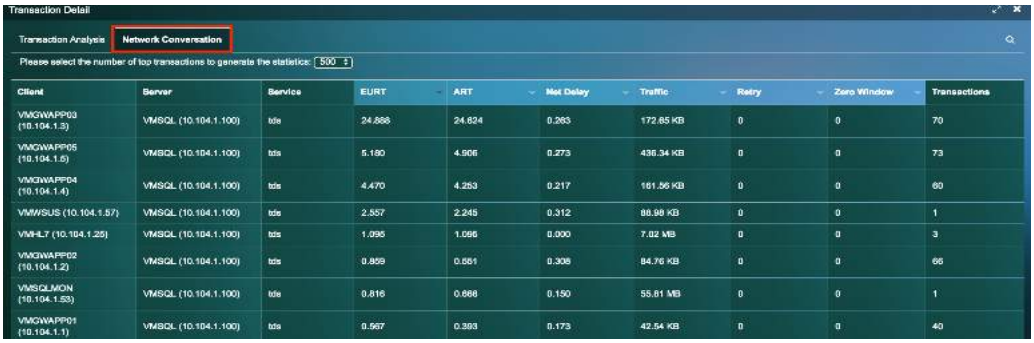

8.3. Service Grouping

Service Grouping page shows a list of all mission critical VM's servicing applications that are essential for the smooth functioning of the datacenter.

8.3.1. Adding a VM to the service resources page

In the service grouping page, the user can create groups and link VM's that belong in those groups.

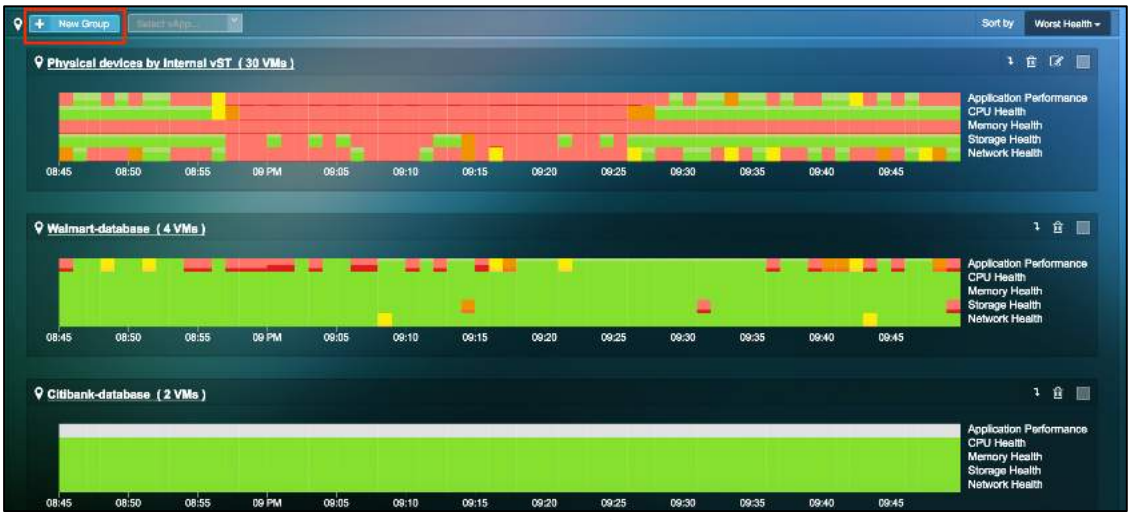

Figure 8.19: Create a new group in the Service Grouping page

VM's that are co-dependent must be added to the group. There are multiple ways to add VM's into Service Groups. The easiest way from the dashboard is to click on the virtual machine of interest, and "Add to Service group".

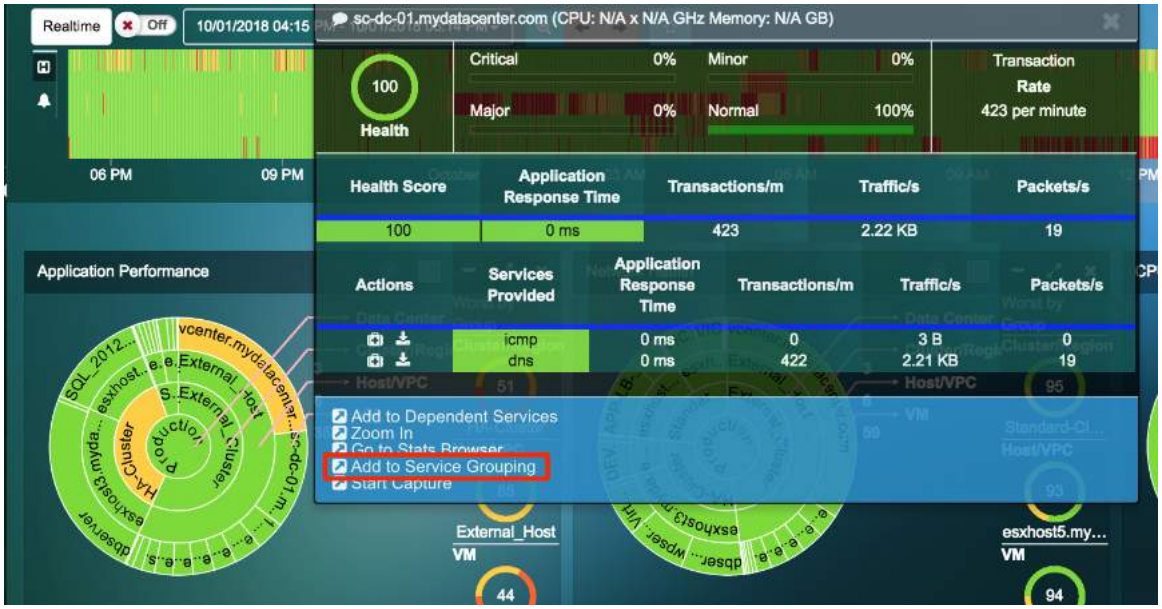

Figure 8.20: Add VM to critical resource

Add the VM to the correct group in to view it from the service grouping page.

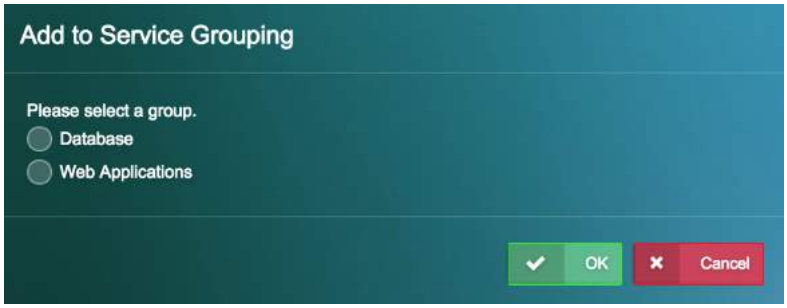

Figure 8.21: Critical resource group

Additionally, multiple VM's can be added to the critical resource view through the application dependency maps, under the application analysis view.

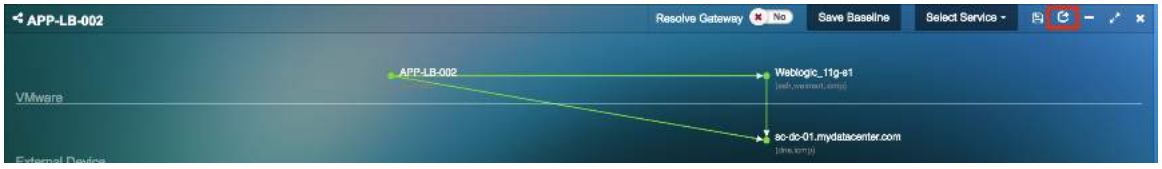

Figure 8.22: Add to service group from dependent service map

When adding VMs into service groups for Application Dependency Mapping, more parameters like the number of transactions and traffic volume can be used to better determine which VMs to add. In this new release we have also enhanced the filtering by VM and service capabilities, by allowing multi-selection for service type.

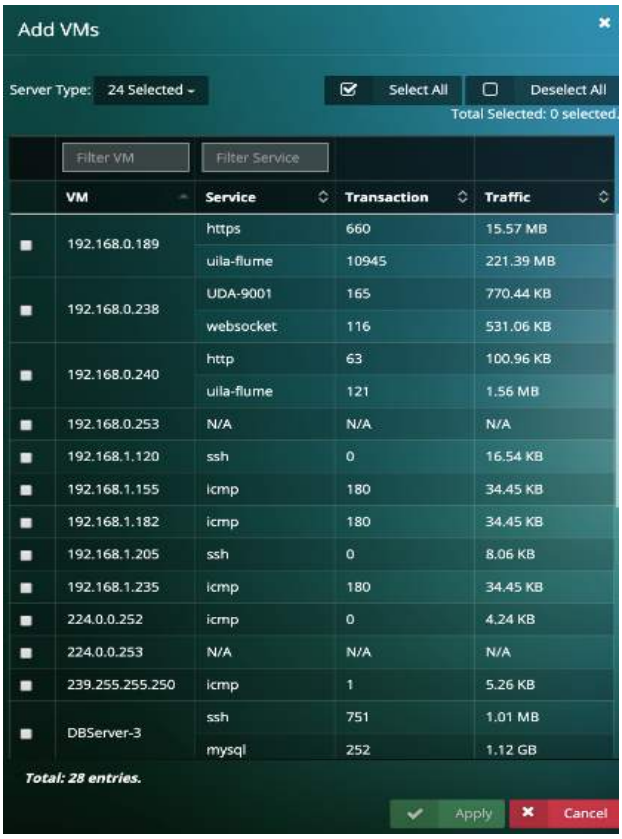

Figure 8.23: Filtering options while adding VMs

You can also use a variety of options from the Actions button. A VDI use case is that users want to see all VDI clients hitting certain backend servers without manually editing the service group.

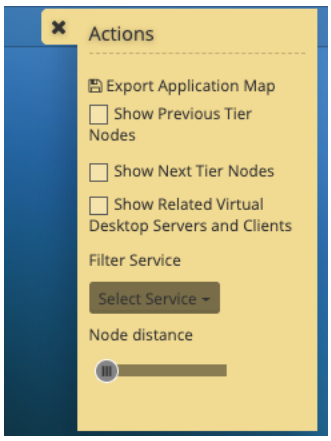

Figure 8.24: Actions option

8.3.2. Monitoring a Service Group

On the critical resource page, click on the group that needs to be monitored.

Figure 8.23: Service groups

The group will show the topology and the list of virtual machines that belong to the group.

Figure 8.24: Service groups

The virtual machine metrics can be seen by clicking on the '+' and expanding the VM.

You can build-out your Application Dependency Maps on a tier-by-tier basis, to provide you with the ability to visualize dependencies that matter to you. This editing capability allows you to visualize dependent servers as well as clients. This can be added by selecting any VM and then choosing the Add Dependent Server or Client option This feature is only available in the Service Grouping section of the application.

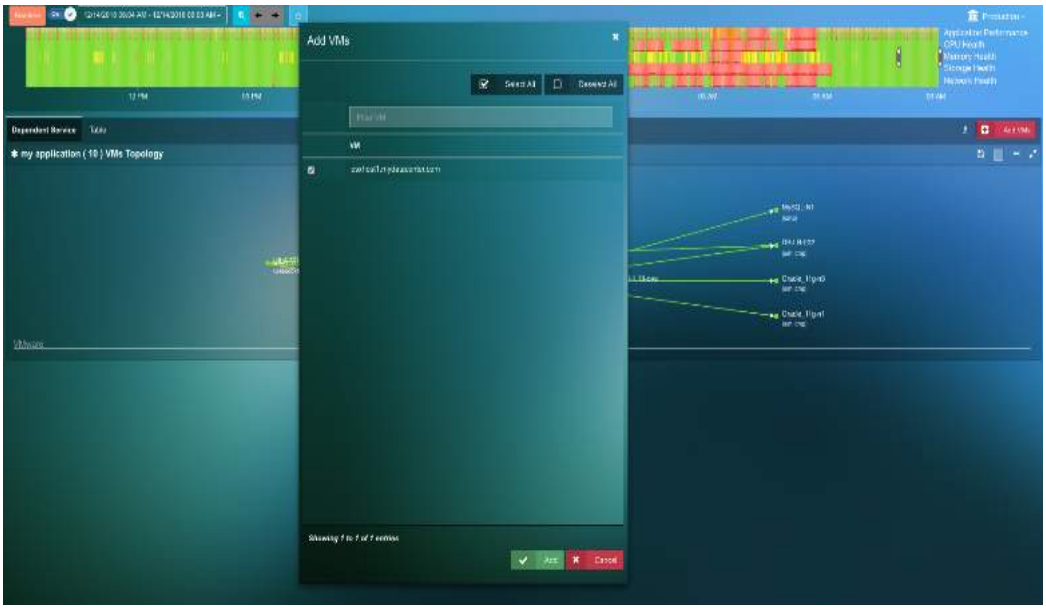

8.3.3.Conversation Map

Users can visualize the applications or services in use on the VMs. For example, this can be very helpful to visualize applications in use on the VDI desktops.

8.3.4.Service Groups based on Port Group

Users can create Service Group for a particular port group directly within the Uila system. Just select the pre-populated port group from the drop down list, and you are all set.

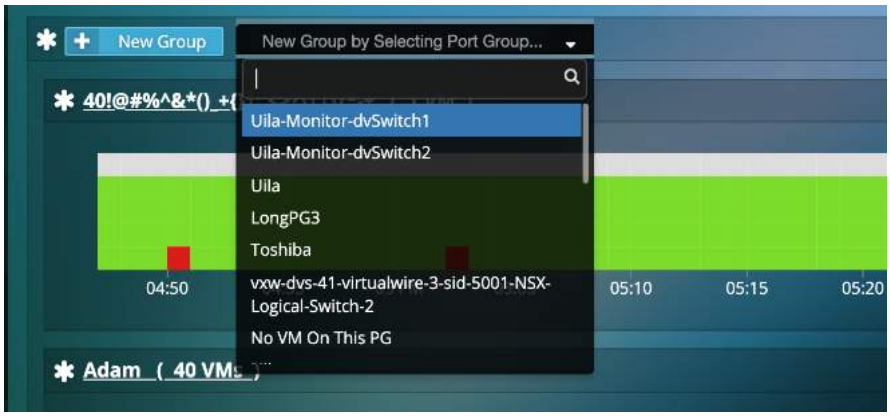

8.4. Service availability

Service availability provides an easy to view interface for mission critical services running in the user's environment. It provides the status of the service along the with the uptime. This feature would be used to ensure all systems and ports of a critical VM are up and functional. If any of the services or VM go down, the user will be able to identify the root cause quickly.

Figure 8.25: Server availability view

8.4.1. Add to Service availability view

Services can be added to the server availability through the critical resources view.

Figure 8.26: Add service to critical resources

Once clicked, input the IP address and the port number to be monitored.

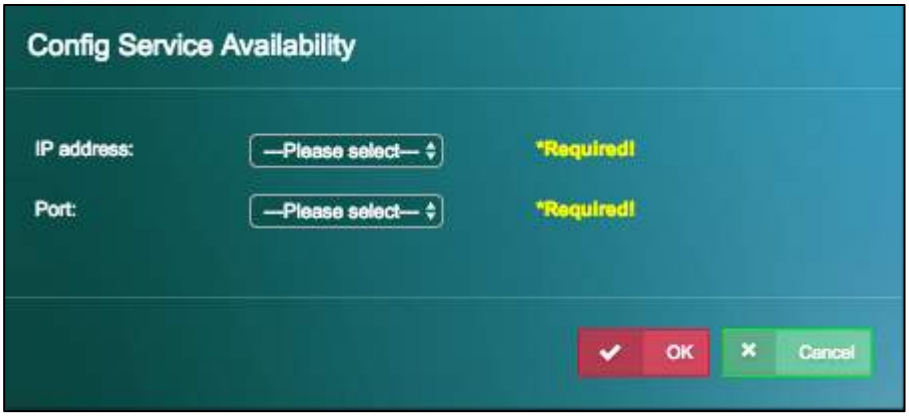

Figure 8.27: Service availability configuration

8.5. End User Experience

Uila measures end user experience for remote sites as well as servers with mission critical functionalities. The user experience is calculated as the sum of application response time, data delivery time and network delay time.

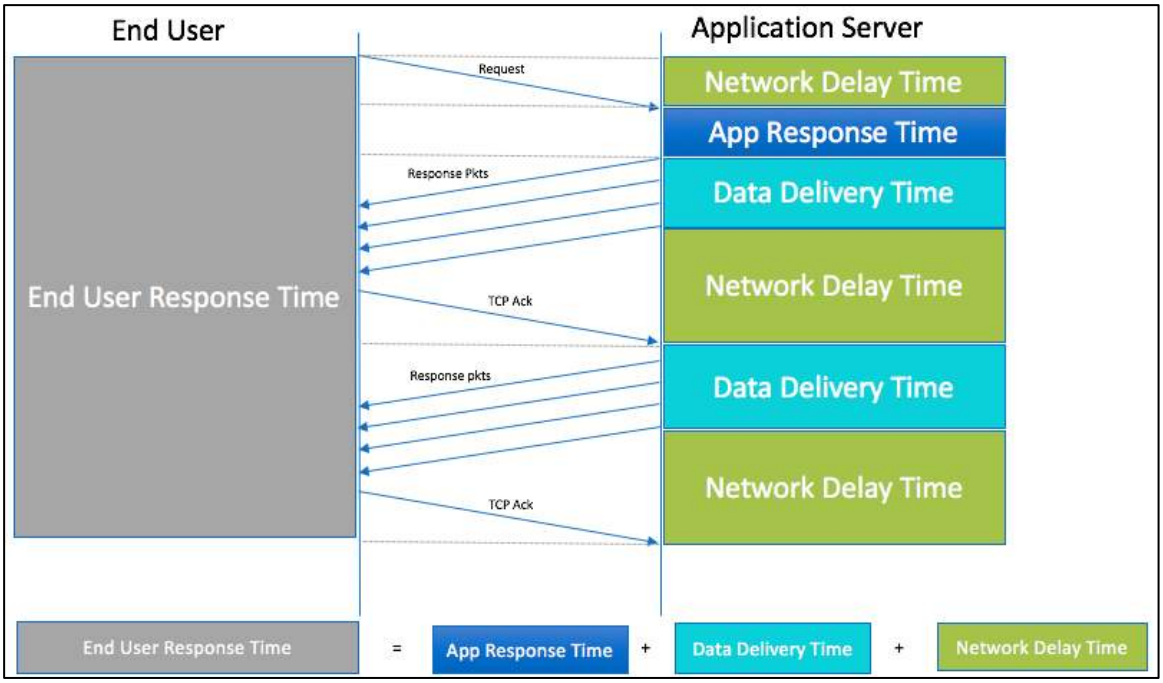

Figure 8.28: End user response calculation

Utilizing the end user experience, the user can identify where the performance issues lie and pin-point them to either server or the underlying network based on the color coding as shown in the Table 16.1. On this page, you can visualize the timeline based on health, Application Response Time or Traffic.

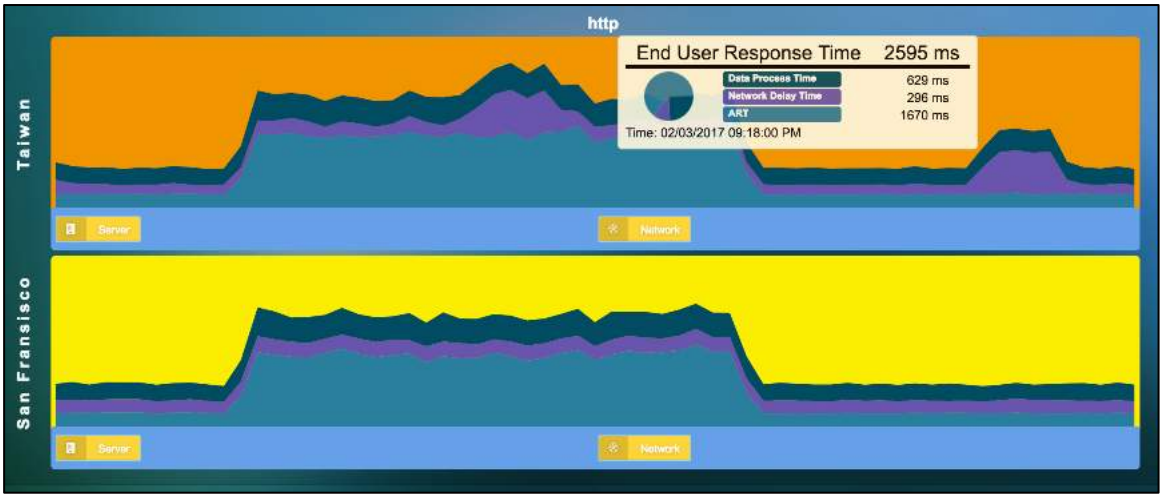

Figure 8.29: End user response time broken down into data process, ART and network delay time.

Component	Normal (Green)	Minor (Yellow)	Major (Orange)	Critical (Red)
Server	Less than 5% from baseline	5-10% from baseline	10-20% from baseline	Over 20% from baseline
Network	Less than 5% from baseline	5-10% from baseline	10-20% from baseline	Over 20% from baseline

Block	Less than 5% from baseline	5-10% from baseline	10-20% from baseline	Over 20% from baseline

Table 8.6: Color codes for User experience

8.5.1.Slow end user response time due to application server

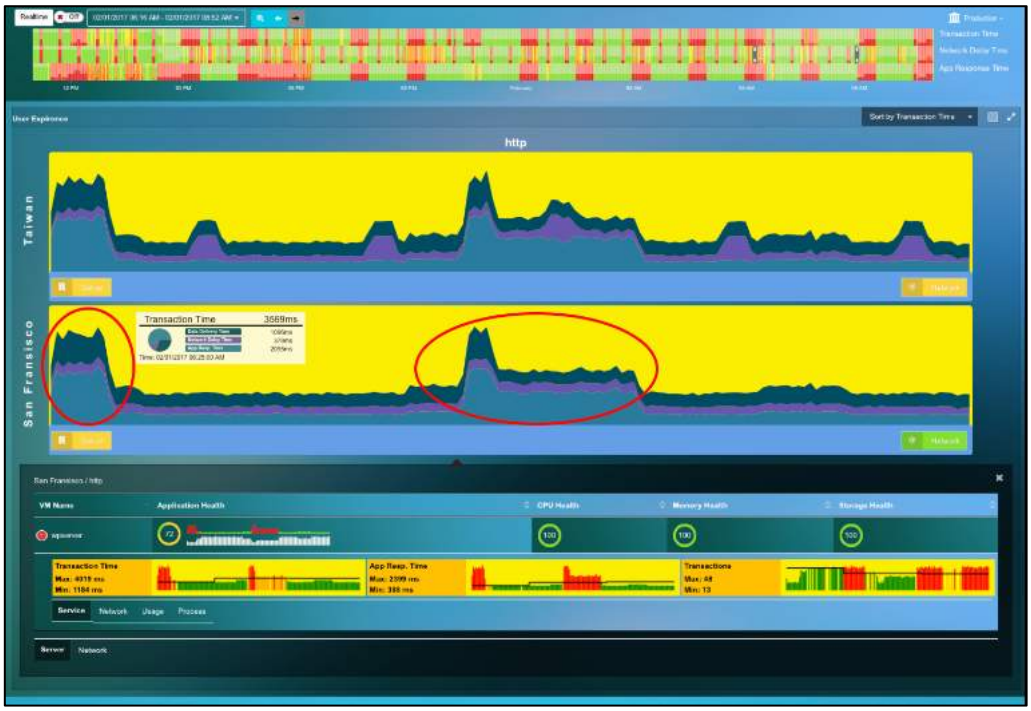

Figure 8.30: Slow end user response time due to application server

To get detailed information regarding application server performance, click on “Server”. The virtual machines hosted on the server will show up and click on the VM that is of concern based on the CPU, memory and storage health.

The end user experience page allows the user to identify the dependent services and get to the root cause of an application slow down.

Figure 8.31: Dependent services within end user response page

By clicking on the deteriorated service, Uila will show up the root cause analysis page with the correlated root cause with CPU, Memory and Storage.

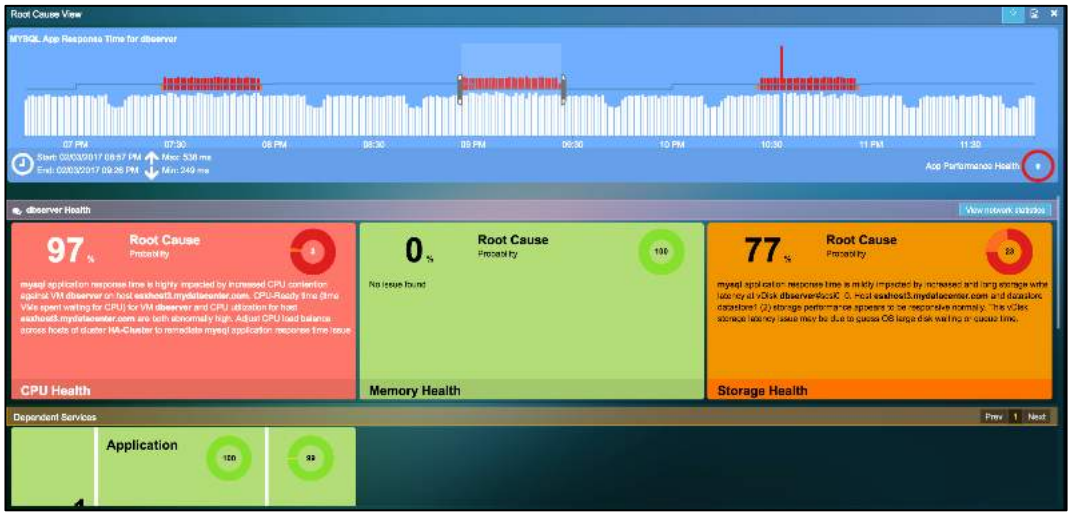

Figure 8.32: Root cause view

8.5.2.Slow end user response time due to Network

As seen in the Fig 15.5(below) we can click on “Network” to understand issues between the remote site and the host. Detailed information such as Network delay time and retransmissions are provided to further analyze the issue.

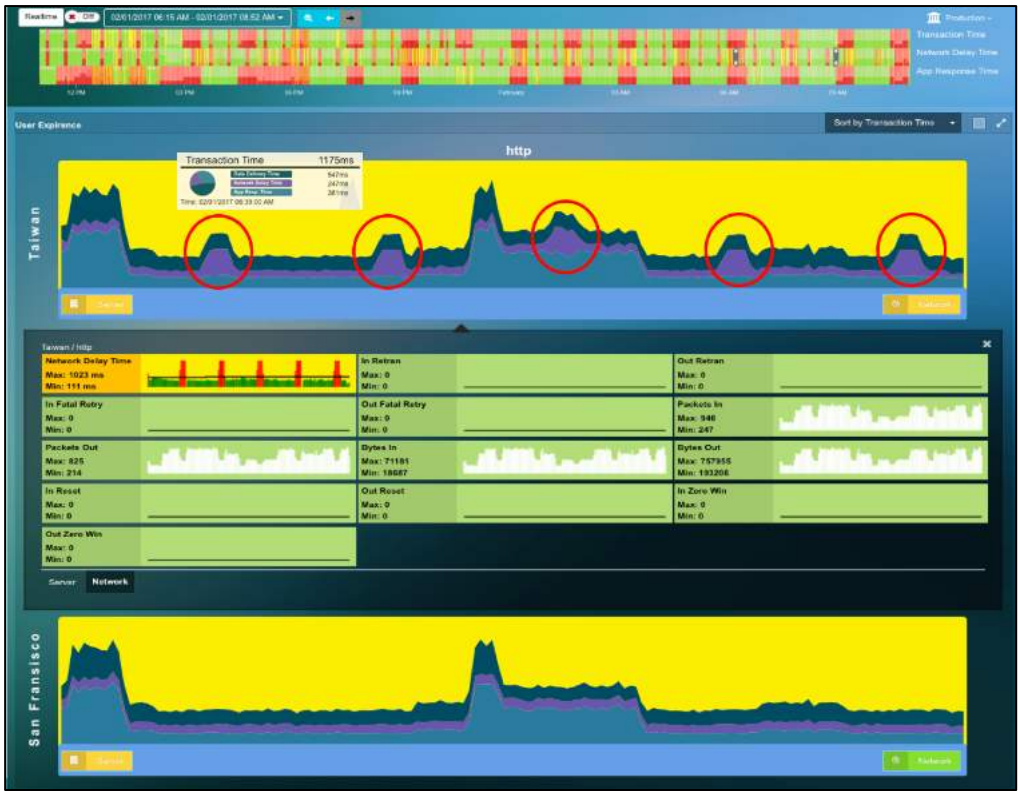

Figure 8.33: Slow end user response due to Network

9. Infrastructure

9.1. Network Analysis

Network Analysis view has a collection of visualization tools; Flow Analysis, Network Conversion, and Table views. Each view is specifically designed to enhance your ability to quickly:

- Identify which infrastructure entities are impacting the Network Health in the Time Frame that is being monitored (one with the Red or Orange color)
- Review network round trip time, application response time and traffic volume of each application service (Classifier) of the respective entity.
- Facilitate further drill down to correlate Application performance impacts.

Network Analysis view is directly launched from the Tool Pane menu, and it consists four tabs (views):

- Flow Analysis view: Visualize how your vAPP network traffic traverses through physical devices (ToR switches, hosts), virtual entities (vSwitch, Port Group, vAPP, VM), and finally, to Application Services (or Classifier) in the data center.
- Network Conversation view: See top-N (100) network traffic volume pairs between VM's and applications served by the VM, and its associated network performance and application performance metrics.
- Network Table view: Organize by all VM's in table view. See Chapter 7.3 Network Performance Metrics
- Alarm View: List of Network alerts generated; Round Trip Time (RTT), Virtual Packet Drops, TCP Fatal Retry, or Reset that exceeds thresholds.

9.1.1. Flow Analysis View

Flow Analysis diagram (also called Sankey diagram) is a powerful visualization tool to show you how your vAPP network traffic are traversing across physical devices (ToR switches, hosts, etc.), virtual entities (vSwitch, Port Group, vAPP, VM), and finally, to Application Services (Classifier) inside your entire data center. You can quickly identify where the network traffic hot spots are, and if they are impacting your application performance. See the sample graphic view below:

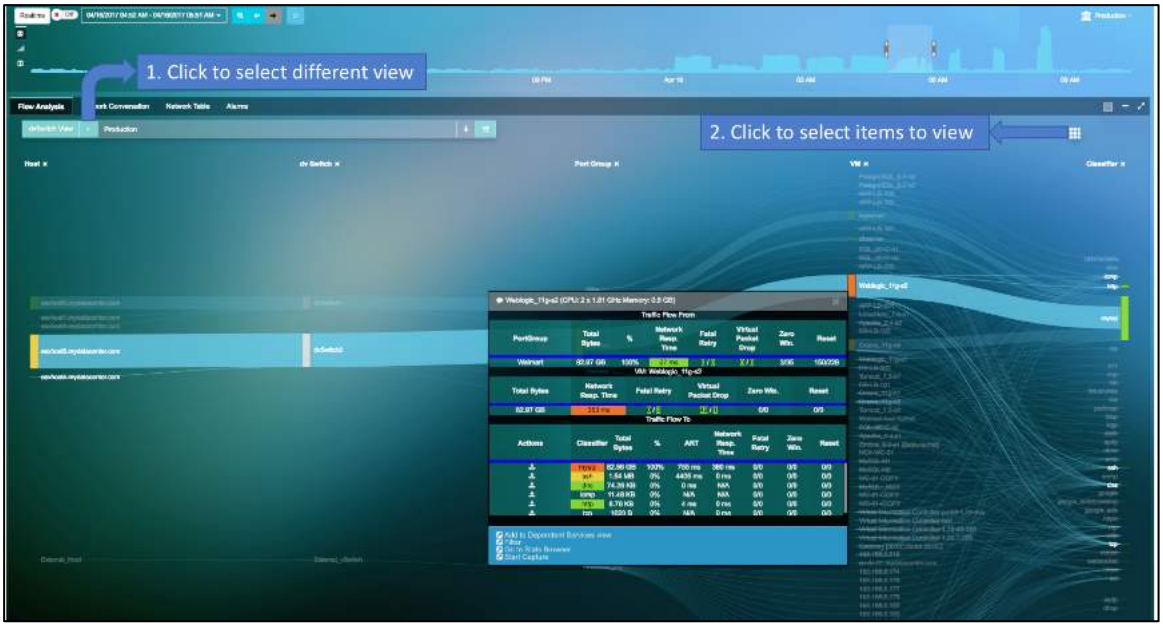

Figure 9.1: Flow Analysis View

Additional Drop-Down list and Buttons in Fig 10.1:

1. Click to display a drop-down list to select a specific view of:

Figure 9.2: Flow Analysis View

2. Click  to display a selection box to select which infrastructure components to display

Figure 9.3: Flow Analysis View

- Select the entities that you wish to display in the Flow Analysis diagram.

Graphic	Definition	Mouse Over Information	Click Action
	Name of physical or virtual entity. Color reflects the network round trip time grading at this entity.	Review network round trip time, application response time and traffic volume of each application service (Classifier) of the respective entity.	Enable *Analyze Application Performance*. Launch Application Topology with filtered view.

Table 9.1: Flow Analysis Graphic

9.1.2. Network Conversation View

Network Conversation provides three types of diagrams to view network traffic volume pairs between VM's and applications served by the VM, and the associated network performance and application performance metrics

- **Top-N Chord View** -

Top-N Chord view displays the top 100 highest network traffic volume VM pairs.

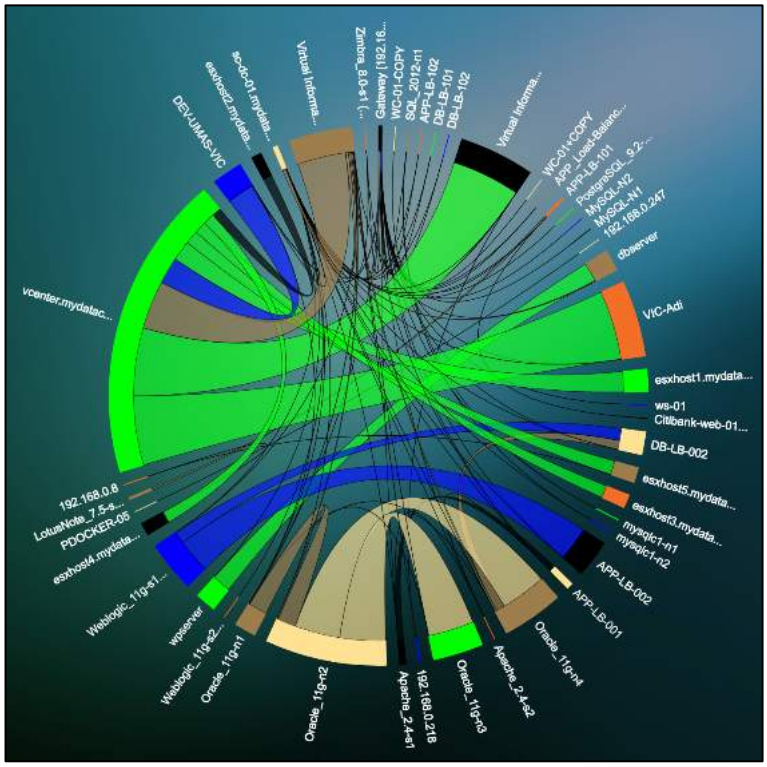

Figure 9.4: Top-N Chord View

- **Top-N Sankey View -**

Top-N Sankey view displays the top 100 highest network traffic volume VM pairs from left to right.

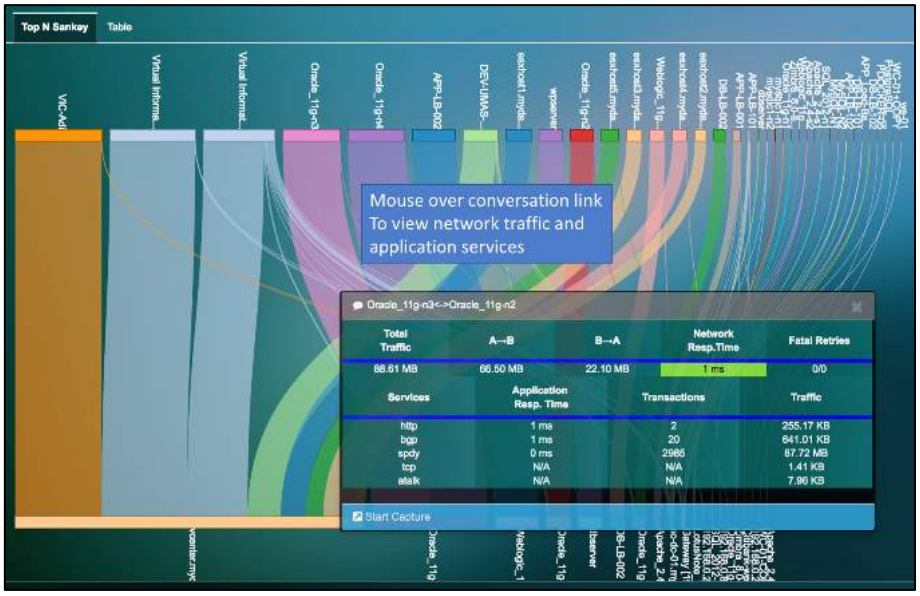

Figure 9.5: Top-N Sankey View

9.1.3. Network Alarm View

Network Alarm view displays network alerts when network performance metrics are above the baseline thresholds. See Chapter 7.3 Network Performance Metric and Chapter 5.2 Health Score and Alarm Definitions.

Network Alarm view provides a detail list of what performances metrics that cause each network alert in the time matrix window you selected. Expand the time matrix window will show more alerts (if any) that were generated in the expanded time slot. If any application service shows performance issue, the name the application service will be displayed in the 'Services'. However, both the network alert and the application performance issues exhibit at the same time do not imply that the cause of application slow is related to networking issue. You need to select and click the root cause view to find the actual root cause.

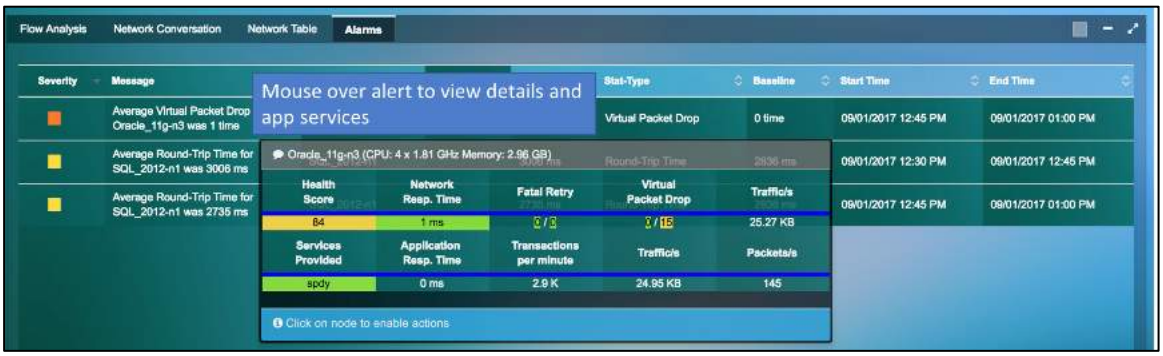

Figure 9.6: Network alarms table

9.2. Network Device Monitoring

Uila users can pinpoint the performance bottleneck down to the network for any dependency chain for a multi-tier application. Users are armed with operational insights on network devices, such as switches, routers, load balancers, firewalls, etc. with detailed info into the availability status, utilization, congestion, errors, discards. In

addition, users get full visibility into connected VMs for every single network switch port and its respective Application, CPU, Memory and Storage Health to pinpoint performance challenges due to the network device bottleneck. For remote location monitoring, in addition to its existing end-user experience monitoring capability to measure the performance from the end-user's perspective & proactively identify issues, users can visualize the status of the WAN link and the interconnection status with the rest of the switch fabric.

The Network Device view will display all the network devices (switches, routers, firewalls, load balancers) along with their port information in the main window pane. For each network device, you can obtain detailed status and configuration settings for network devices including vendor, model, OS versions, uptime, serial number, VTP domain, detailed description, IP/MAC address, etc.

Open (no Ethernet cable plugged in) and down/disabled ports are indicated by a "Hollow" port icon. If it is "green", the port is open, while "red" indicates that the port is down or disabled.

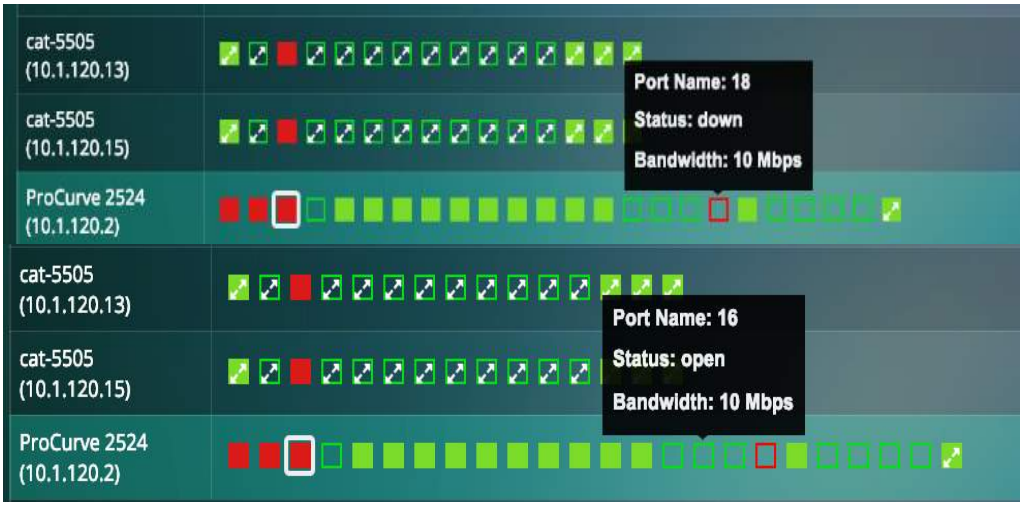

For each port, you can visualize the following statistics in a chart format (1st tab).

- In/Out Utilization
- In/Out Discards
- In/Out Errors
- In/Out Unicast Packets
- In/Out Non Unicast Packets
- In/Out Octets
- Queue Length
- Unknown Protocol packets

The following charts define the solid colors seen for the ports in the User Interface.

- In/Out Utilization
- In/Out Discards
- In/Out Errors

You can set the thresholds for the parameters from the “Threshold Settings” tab for individual ports.

The Default baselines are as follows:

- Utilization: 80%
- Discards: 10,000 pkts/min
- Errors: 100 pkts/min

Alarm is generated based on the performance metric’s delta from the baseline. Alarm is generated every 15 minutes by default.

Threshold is defined as the % value that crosses the baseline.

Severity is a user definable indicator to help identify the criticality of the performance metrics monitored to alert user if an entity or entities is (are) about to impact the Application’s performance.

Delta from Baseline	Alarm Severity	Color
Less or equal to 5%	**Normal**	Green
Between 5% and 10%, including 10%	**Minor (1)**	Yellow
Between 10% and 20%, including 20%	**Major (2)**	Orange
Above 20%	**Critical (3)**	**Red**

Note: These standard color definitions are applied throughout Uila User Interfaces for consistence and ease of recognition.

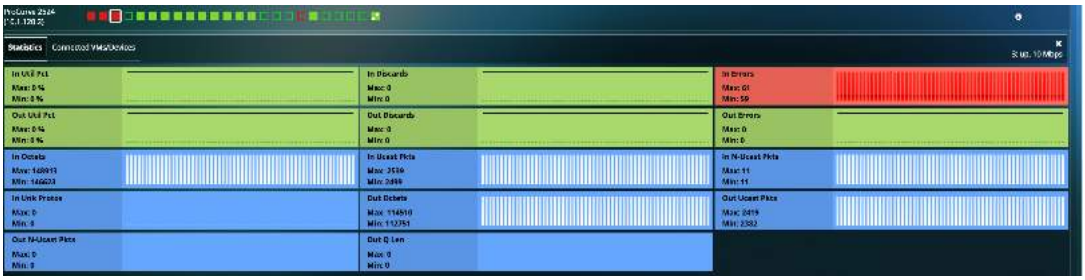

Cross arrow inside the square icon for a port indicates a connection from that port to other switches/routers. The same logic applies for colors, as the solid colors mentioned in question #8. Note: This feature is supported for switches and routers only, and not for other network devices.

This can be used to show the status of the WAN link and the interconnection status with the rest of your switch fabric.

Also for each port, you can visualize the Connected VMs/Devices in the next tab. For every VM, you can visualize the Application, Network, CPU, Memory and Storage Health. Further VM statistics (Usage, Alarms, Process, Dependent Services, etc.) can be obtained by clicking on the VM name.

You can also visualize alarms in the alarm tab within Network Device if a particular port is congested (high utilization) or has errors (errors, discards).

9.3. CPU Analysis

CPU Analysis view has a collection of visualization tools; Circle Packing, Tree, Table and Alarm views, each is specifically designed to enhance your ability to quickly:

- Identify the infrastructure entities impacting the CPU Health in the Time Frame that is being monitored (one with the Red or Orange color)
- Review application response time and traffic volume of each application service (Classifier) related to CPU usage %, CPU MHz and CPU ready % with respect to each element.
- Facilitate further drill down to correlate Application performance impacts by CPU performance.

CPU Analysis view is directly launched from the Tool Pane menu, and it consists four tabs (views):

- Circle Packing view: Visualize CPU Capacity, and CPU Usage of each VM, host and cluster within the data center infrastructure. The size of the circle is determined by CPU capacity, while the width of the ring is related to the CPU usage of each element.

- Tree view: Alternative view to visualize CPU Capacity, and CPU Usage of each VM, host and cluster within the data center infrastructure. The size of the circle is determined by CPU capacity, while the width of the ribbon is related to the CPU usage of each element.
- Table view: Organize in table view to sort by performance grade of the VM. Refer to Chapter 7.5 CPU Performance Metrics for details.

- Alarm View: List of CPU alerts generated; CPU Usage %, or CPU Ready time (in %) that exceeds thresholds.

9.3.1. Circle Packing View

Circle Packing view allows you to visualize CPU capacity, and CPU usage, and Health of each VM, host and cluster within the data center infrastructure. The size of the circle is determined by CPU capacity, while the width of the ring is related to the CPU usage of each element. When CPU usage percentage reaches certain thresholds, the circle turns yellow, orange or red, indicating which entity is busy. By comparing the size of all VM circles under a host, you can quickly know how evenly a CPU capacities (vCPU cores) are allocated across all VMs. Sometimes, a big VM in term of CPU core numbers may impact its peer VM's performance. Mouse over the element that exhibits health performance issue, you can further drill down to reveal how application response time is impacted.

Figure 9.7: CPU Circle packing view

9.3.2. Tree View

Tree view is an alternative view to allow you to visualize CPU capacity, and CPU usage, and Health of each VM, host and cluster within the data center infrastructure. The size of the circle is determined by CPU capacity, while the width of the ribbon (same as the size of the pie slice) is related to the CPU usage of each element. When CPU usage percentage reaches certain thresholds, a circle turns yellow, orange or red, indicating which entity is busy. By comparing the
size of all VM circles under a host, you can quickly know how evenly a CPU capacity (vCPU cores) are allocated across all VMs. Sometimes, a big VM in term of CPU core numbers may impact its peer VM's performance. Mouse over the element that exhibits health performance issue, you can further drill down to reveal how the application response time is impacted.

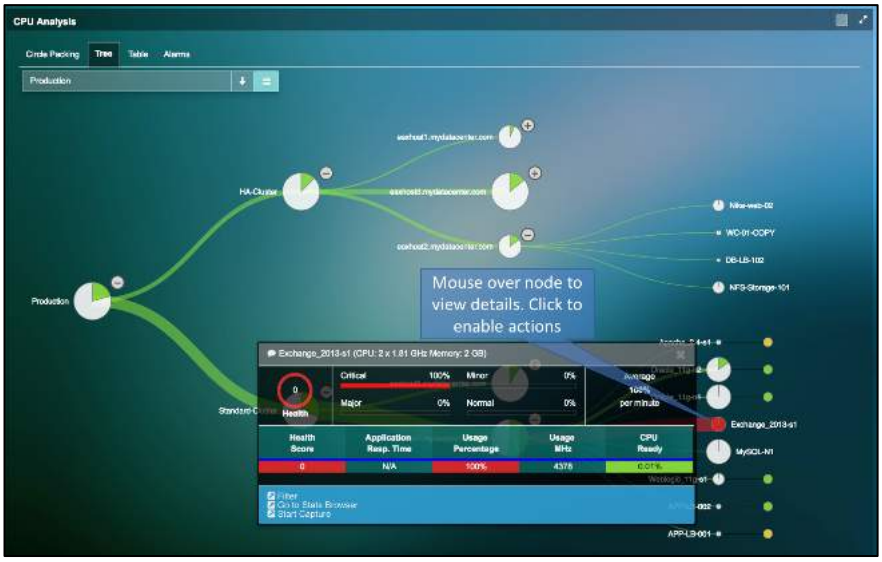

Figure 9.8: CPU tree view

9.3.3. Alarm View

CPU Alarm view displays CPU performance alerts when CPU usage or CPU ready metric is above the baseline thresholds. See Chapter 7.5 CPU Performance Metric and Chapter 5.2 Health Score and Alarm Definitions.

CPU Alarm view provides a detail list of what performances metrics that cause each CPU performance alert in the time matrix window you selected. Expand the time matrix window will show more alerts (if any) that were generated in the expanded time slot. If any application service shows performance issue, the name the application service will be displayed in the 'Services' column. However, if both the CPU alert and the application performance issues exhibit at the same time, it does not imply that the cause of application slowness is related to CPU issue. You need to select and click the root cause view to find the actual root cause(s).

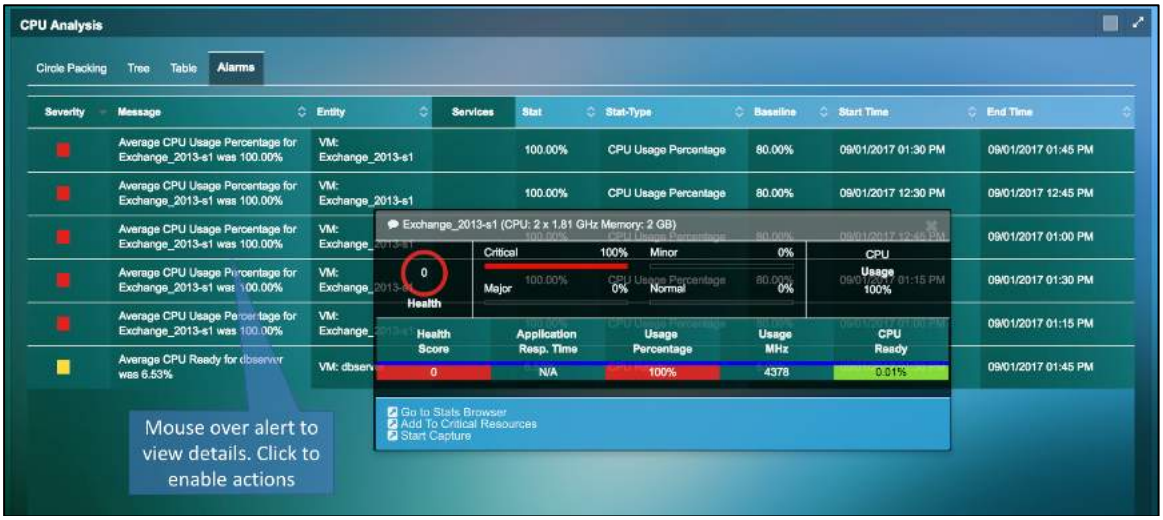

Figure 9.9: CPU alarm view

9.4. Memory Analysis

Memory Analysis view has a collection of visualization tools; Circle Packing, Tree, Table and Alarm views, each is specifically designed to enhance your ability to quickly:

- Identify which infrastructure entities are impacting the Memory Health in the Time Frame that is being monitored (one with the Red or Orange color)
- Review application response time and traffic volume of each application service (Classifier) related to Memory usage %, and CPU Swap Wait time with respect to each element.
- Facilitate further drill down to correlate Application performance impacted by Memory performance.

Memory Analysis view is directly launched from the Tool Pane menu, and it consists four tabs (views):

- Circle Packing view: Visualize Memory Capacity, and Memory Usage of each VM, host and cluster within the data center infrastructure. The size of the circle is determined by Memory capacity allocated, while the width of the ring is related to the Memory usage of each element.
- Tree view: Alternative view to visualize Memory Capacity, and Memory Usage of each VM, host and cluster within the data center infrastructure. The size of the circle is determined by Memory capacity, while the width of the ribbon is related to the Memory usage of each element.
- Table view: Organize in table view to sort by performance grade of the VM. Refer to Chapter 7.6 Memory Performance Metrics for details.
- Alarm View: List of Memory alerts generated; Memory Usage %, or CPU Swap Wait time that exceeds thresholds.

9.4.1. Circle Packing View

Circle Packing view allows you to visualize Memory capacity, Memory usage, and Health of each VM, host and cluster within the data center infrastructure. The size of the circle is determined by Memory capacity, while the width of the ring is related to the Memory usage of each element. When Memory usage percentage reaches certain thresholds, a circle turns yellow, orange or red, indicating which entity is busy. By comparing the size of all VM circles under a host, you can quickly know how evenly a Memory capacity are allocated across all VMs. Sometimes, a high Memory usage VM may require allocation of more memory compared to VM's that are less frequently run. Mouse over the element that exhibits health performance issue, you can further drill down to reveal how application response time is impacted.

Figure 9.10: Memory circle packing view

9.4.2.Tree View

Tree view is an alternative view to allow you to visualize Memory capacity, and Memory usage, and Health of each VM, host and cluster within the data center infrastructure. The size of the circle is determined by Memory capacity, while the width of the ribbon (same as the size of the pie slice) is related to the Memory usage of each element. When Memory usage % reaches certain thresholds, a circle turns yellow, orange or red, indicating which entity is busy. By comparing the size of all VM circles under a host, you can quickly know how evenly a Memory capacity are allocated across all VMs. Mouse over the element that exhibits health performance issue, you can further drill down to reveal how application response time is impacted.

Figure 9.11: Memory tree view

9.4.3.Alarm View

Memory Alarm view displays Memory performance alerts when Memory usage or CPU Swap Wait time metric is above the baseline thresholds. See Chapter 7.6 Memory Performance Metric and Chapter 5.2 Health Score and Alarm Definitions.

Memory Alarm view provides a detail list of performances metrics that cause Memory performance alert in the time matrix window that has been selected. Expand the time matrix window will show more alerts (if any) that were generated in the expanded time slot. If any application service shows performance issue, the name the application service will be displayed in the 'Services' column. However, if both the Memory alert and the application performance issues exhibit at the same time, it does not imply that the cause of application slowness is related to Memory issue. You need to select and click the root cause view to further pinpoint the actual root cause(s).

Figure 9.12: Memory alarms view

9.5. Storage Usage

Storage Usage diagram is a visualization tool to show you Storage usage and Health Score within your data center physical or virtual entities. Storage Usage view can be launched from Dashboard's Storage Health color wheel, or directly from the Tool Pane menu.

The figure below shows the navigation method and tool tips in the Storage Usage view.

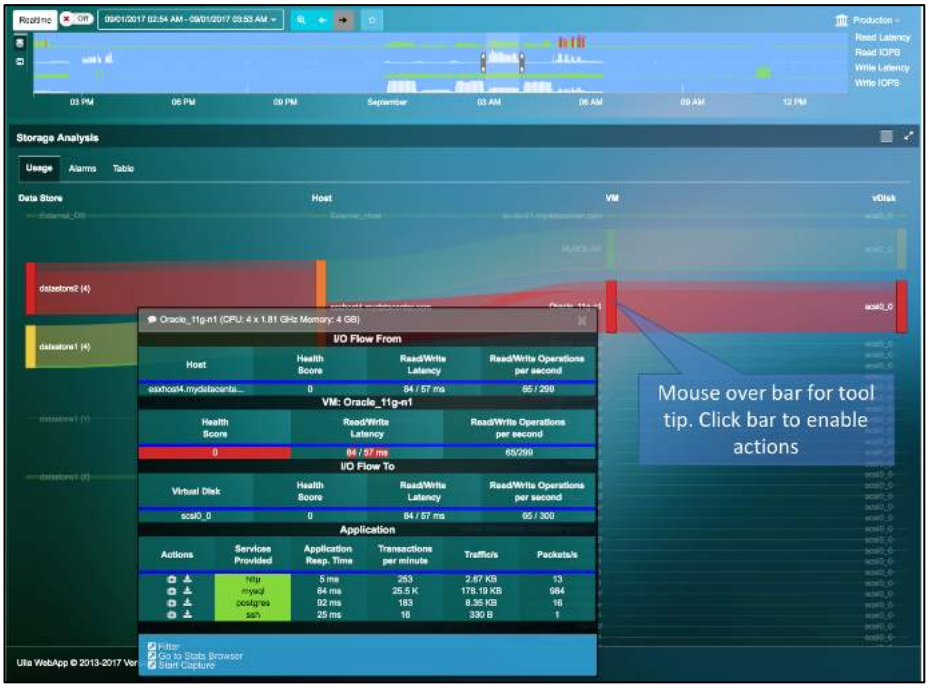

Figure 9.13: Storage Usage View

Refer to Section 7.4, Storage Performance Metric for Storage metric definition, and how metrics are calculated to determine health score and the associated base line values.

To help investigate performance issues, you can place the mouse over the vertical bar of each storage infrastructure component to reveal the health and performance summary of its upstream and downstream neighbors in a Tool Tip.

10. Security

Uila leverages its Deep Packet Inspection (DPI) capability to make use of network packet data as the root of truth and identifies advanced threats that are moving laterally (insider threats). Users can detect and manage cyber alerts and anomalous deviations in dependencies for applications that are business critical to the enterprise organization to bring an unique Application-centric view to cyber threat monitoring. Uila provides the necessary Intelligence & Diligence to reduce the attack surface and becomes a force multiplier for security operations teams. Security and Network teams are automatically alerted to the latest malicious threats and attacks, including malware, exploit kits, outbound traffic issues, C & C threats, etc. In addition to the latest threats, IT teams can confidently track the chain-of-evidence for critical Network and Application workload characteristics in real time to identify anomalous outliers such as dependency changes between the critical application and infrastructure resources, deletion or addition of new VMs, etc.

The time slider for security will indicate the levels of threats that have been identified in the deployment.

Figure 10.1: Security Time slider

10.1 Application Anomaly

You can now visualize Application deviations for your multi-tier applications (created based on Service Groups) indicating anomalous behavior in a single view. In addition to insights into detailed cyber threat event information and outbound traffic behavior to the Internet for the group, you can visualize deviations after the creation of your desired baseline for the application or service. Deviations include unauthorized dependency changes, new applications/services/protocols running on the VMs, additions of unauthorized VMs or tearing down of your mission critical VMs, etc. You can visualize those deviations in the Application Dependency Map and add deviations to the baseline or security policy.

All Service Groups that have been created will appear automatically on this screen. For every service group, Uila will list if there is any deviation from the configured baseline, Cyber threats that have been identified as well as Data Exfiltration transactions.

Figure 10.2: Application Anomaly Overview

The first step is to configure the baseline for the known good time period for the Application dependencies.

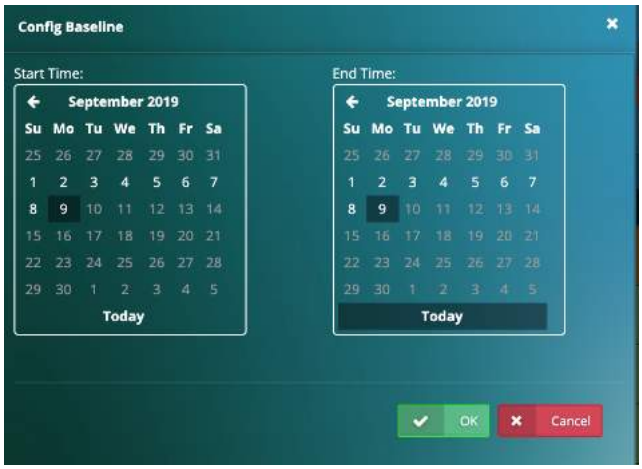

Figure 10.3: Configure Baseline

Once the baseline is configured, users can visualize the application anomalies or deviations in the tabular format, as well as visually as an Application Dependency Map.

Uila will list all the individual deviations taking place for every asset as part of the Service Group. Few examples include, addition or removal of VMs, addition of services, new dependencies, new requests and responses, etc.

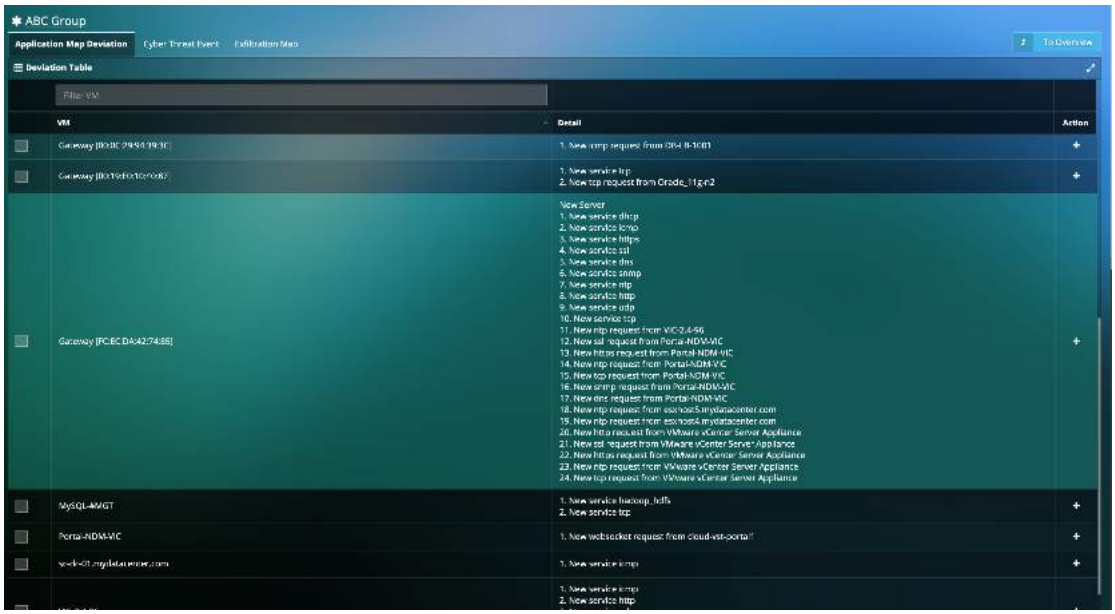

Figure 10.4: Application Anomaly Table

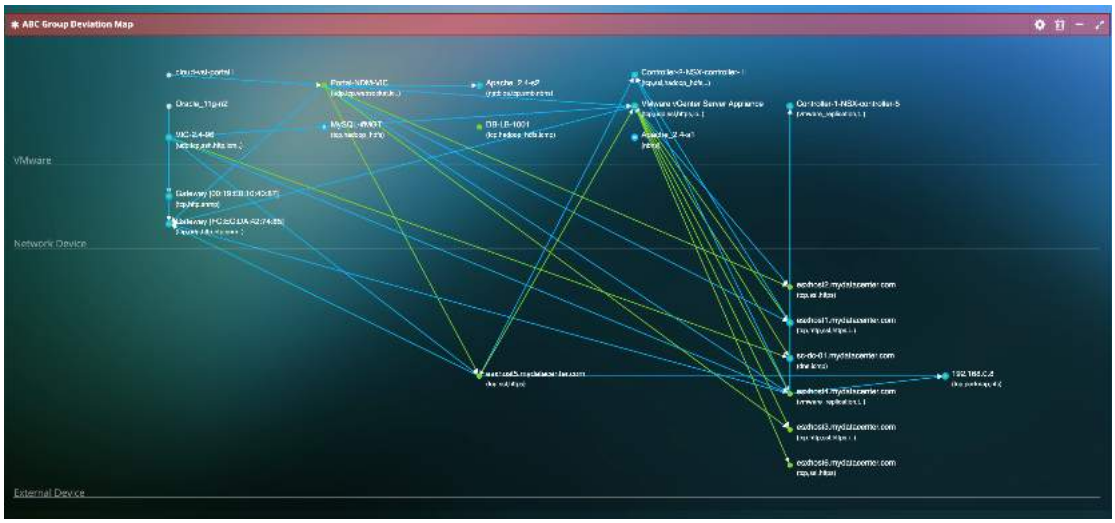

Figure 10.4: Application Anomaly Dependency Map

Users can also visualize the deviations for each individual asset, by checking the individual boxes. Also, if the deviation is expected or valid, it can be directly added to the baseline by clicking the “+” button.

10.2 Cyber Threat Monitoring

Uila users can now get alerted to thousands of cyber threats based on support from the largest group dedicated to advances in the network security industry. These alert categories include malware, exploit kits, port scans, Command and Control threats, OS fingerprinting,

Buffer overflows, SMB probes, Obfuscation, etc. Uila supports latest signature support and updates from the largest group dedicated to advances in the network security industry (Snort, Cisco® Talos Security Intelligence and Research Group, ClamAV). This can be viewed for the entire Data Center or for a Service Group.

Uila provides graphical summary of the following information:

- Threat Severity (Critical, Major or Minor)
- Threat Models or Categories
- Threat Types
- Threat Source and Destination

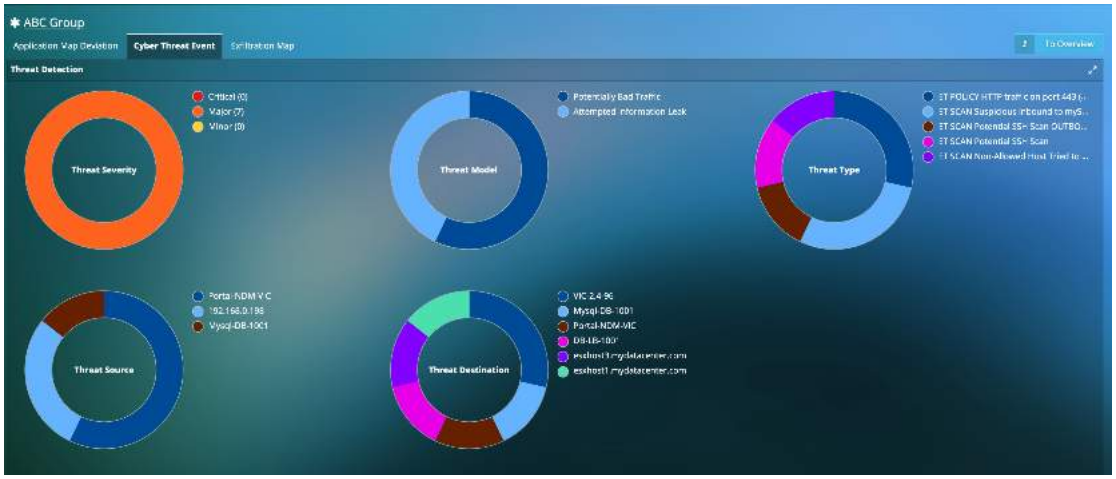

Figure 10.5: Cyber Threat Summary

Each cyber threat is also listed with information on its severity level, threat model, type, source and destination and the event count (tracked on a minute by minute basis).

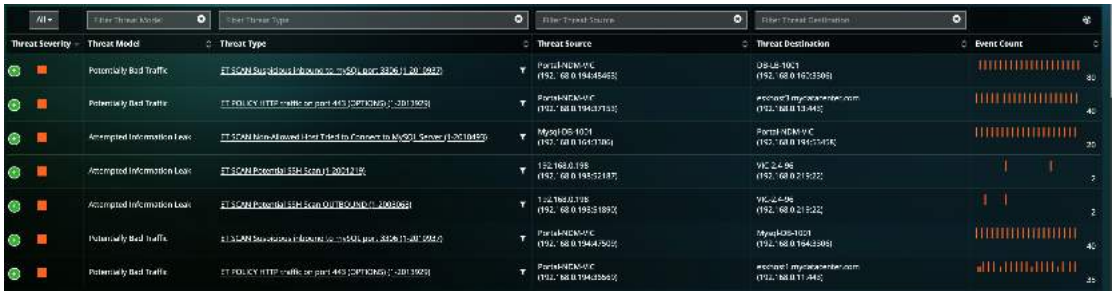

Figure 10.6: Cyber Threat Summary Table

For each of the threats, you are powered with information on the Application Dependencies. Uila highlights the source and destination of the threat (which indicate the attacking or the attacked/compromised entity). As you have visibility into all of the dependencies, you have insights into entities or assets that could get compromised in the future. For example, a webserver that is currently facing an attack, may not be the ultimate goal for the attacker. The goal could be to reach and compromise the database server that is connected to that webserver. Knowing all the dependencies gives you the proactive knowledge into future attacks or vulnerabilities. Also, with Uila you can get access to all transactions at the application level that can be maintained as forensic evidence.

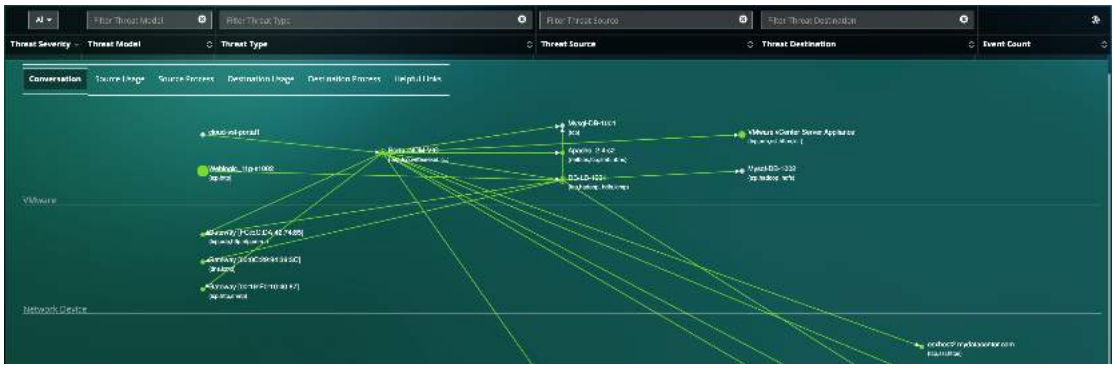

Figure 10.8: Cyber Threat Conversation Maps

You can also apply a variety of display filters to the table to help you focus on cyber threats that matter to you. In the example below, we chose to visualize alerts based on threat models with the term “leak” in it.

Figure 10.8: Cyber Threat Display Filters

For every threat, you have the ability to visualize the impact that the threat has on the entity’s infrastructure (CPU, memory, storage, network stats).

Figure 10.9: Source & Destination Infrastructure usage

You can also visualize the processes running on the source and destination entities

Figure 10.10: Source & Destination Process Information

You can also visualize helpful links on each of the cyber threats. You get expert guidance on those threats, their symptoms, the impact and corrective actions to solve and avoid future reoccurrences.

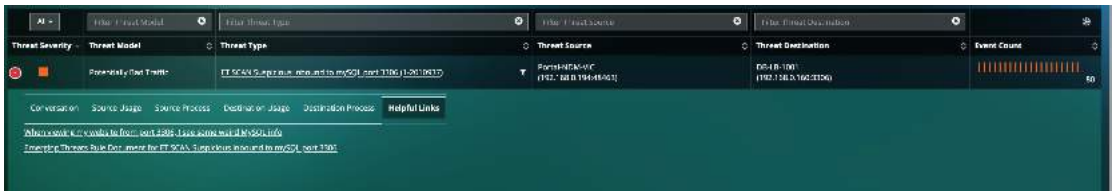

Figure 10.11: Cyber Threat Helpful Links

10.3 Data Exfiltration

Uila users can now map Outbound Traffic from the Data Center to the Internet on a world map to identify and reduce risk associated with general Internet connectivity. You can visualize Outbound traffic details including Internal VM details, Destination IP, Destination Server location, Application/Service for the outbound traffic, etc. This can be viewed for the entire Data Center or for a Service Group.

You also have the option to filter on information that matters to you on this screen as well as the option for visualizing the transactions at the application level and add to dependent services and external devices.

11. Root cause view

The root cause view provides quick root cause analysis of persisting application level issues within the datacenter. The application response time is correlated with the infrastructure (compute, storage and network) as well as the services the problematic VM relies on.

Worst Transaction details are also provided in to help the systems administrator look into the transaction history and troubleshoot the application in case there are no issues on the infrastructure side.

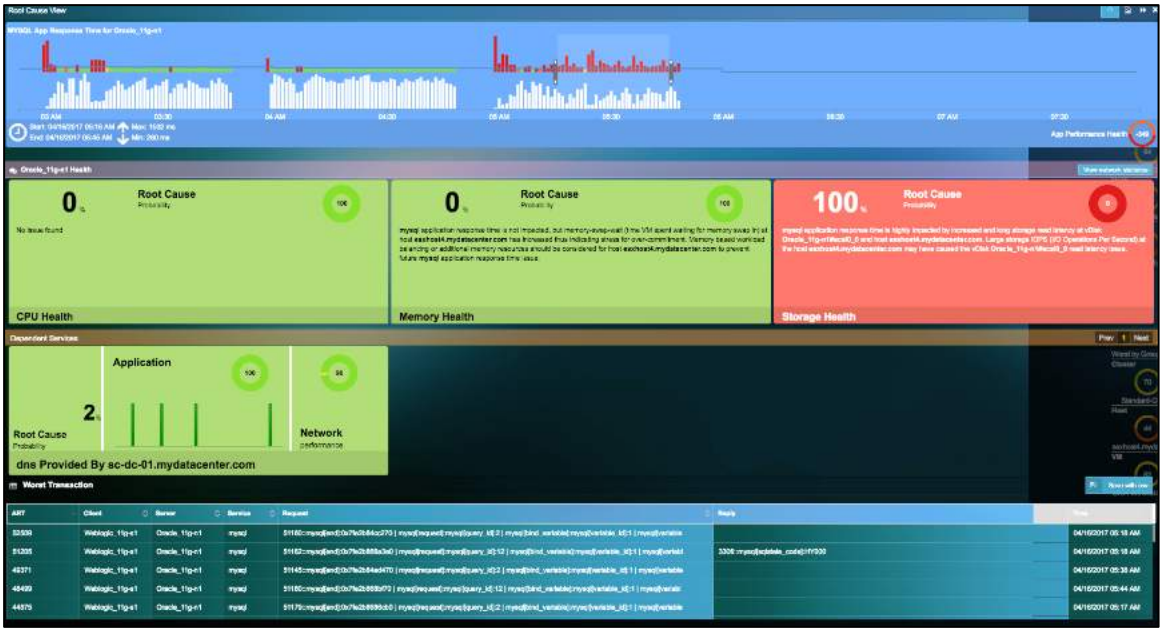

Figure 10.1: Storage Usage View

The user would be able to further drill down by further clicking on the health panes or dependent service pane to get in-depth information.

11.1. CPU Health

Under the CPU health analysis view, Uila can provide detailed information on CPU usage, CPU ready as well as the CPU MHz. This information can help the user analyze the factors responsible for the high ART.

Process level information can also be gathered from the OS through WMI(Windows) or SSH(Linux) integration.

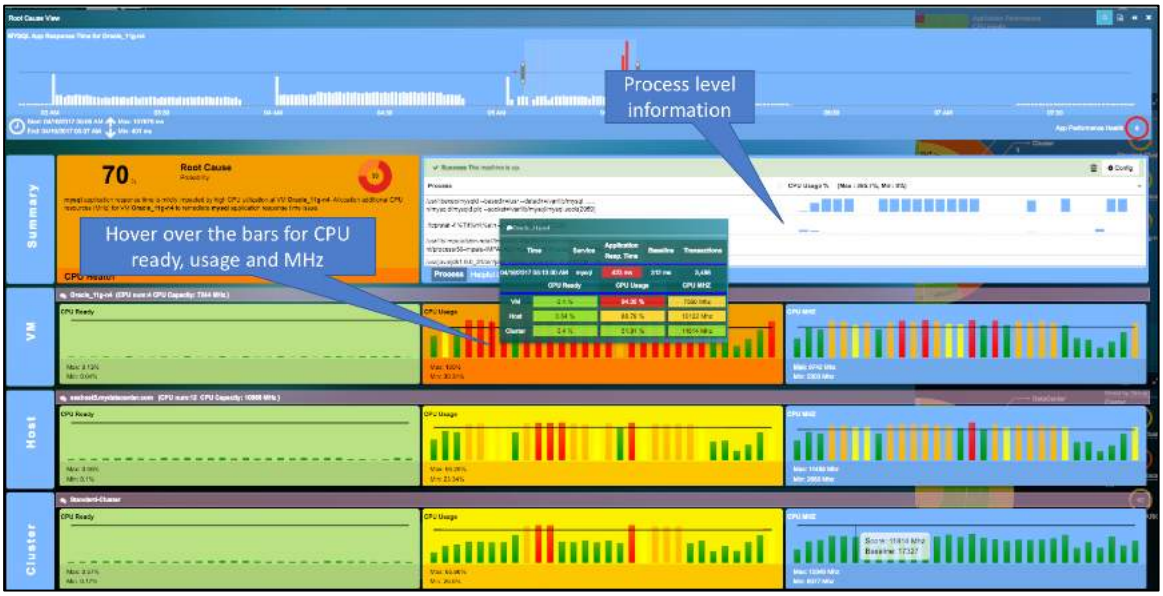

Figure 10.2: CPU Health Root Cause View

11.2. Memory Health

Under the Memory health analysis view, Uila can provide detailed information on Memory usage and CPU Swap wait time. This information can help the user analyze the factors responsible for the high ART.

Process level information can also be gathered from the OS through WMI(Windows) or SSH(Linux) integration.

Figure 10.3: Memory Health Root Cause View

11.3. Storage Health

Under the Storage health analysis view, Uila can provide detailed information on read/write latency and IOPS. This information can help the user analyze the factors responsible for the high ART.

By clicking on the bars, the user can understand the neighboring VM's that share the same resources.

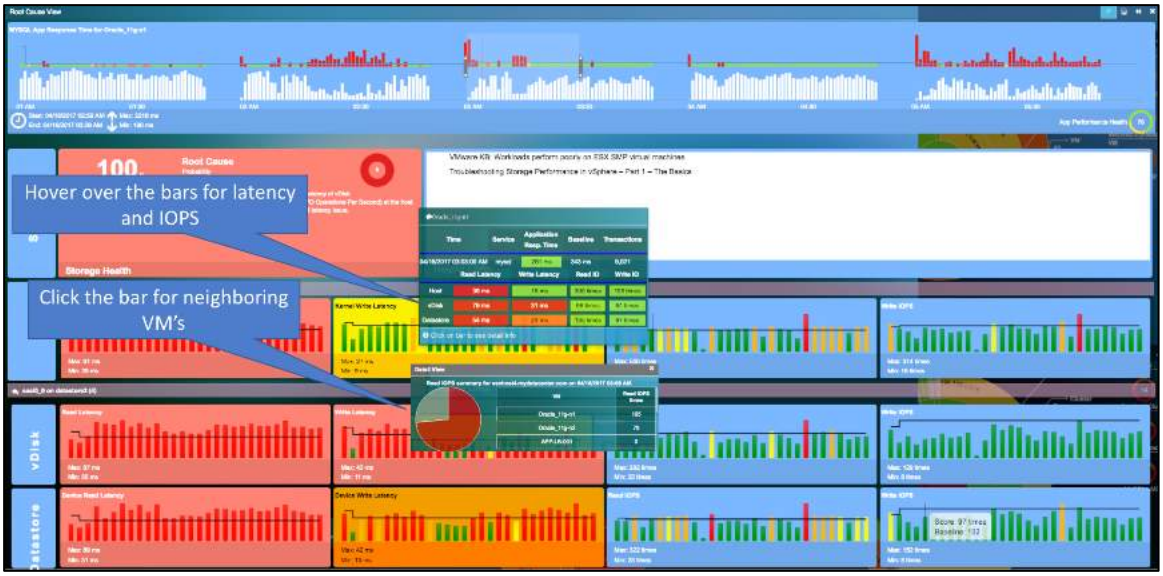

Figure 10.4: Storage Health Root Cause View

12. Stats Browser

Stats Browser is another powerful visualization tool that places all the metrics collected for any of the infrastructure components; Cluster, Host, and VM in one single unified screen view. It is a particularly useful when the root cause of an application performance issue has been identified and the user wishes to further validate it across all the infrastructure metrics.

You also have the option to visualize detailed information that is specific to a server or VM or external IP address. Users are powered with a map that displays all related network, infrastructure and application (service) associated with the VM/Server/IP address. By clicking on any entity in the map, you can then get further details on related metrics and statistics.

The figure below shows the navigation method and tool tips in the Stats Browser view:

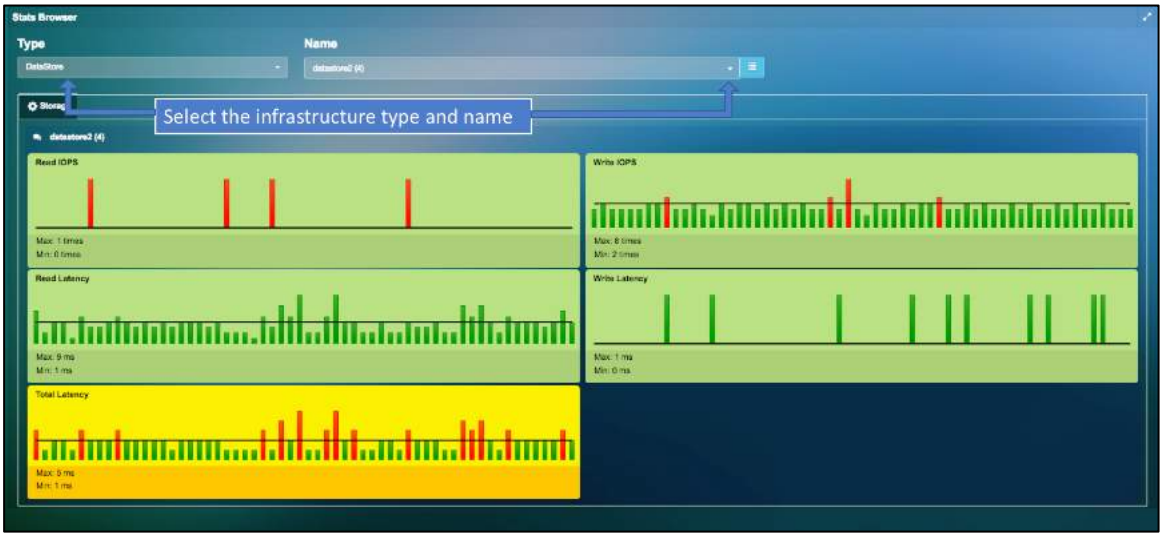

Figure 11.1: Stats browser

Use the Drop-down box below to select Type and name of the specific infrastructure units to view the summary of metrics over time bracket selected:

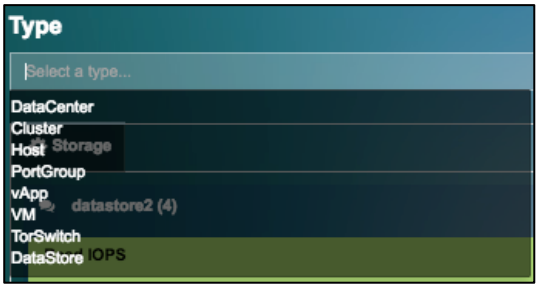

Figure 11.2: Types drop-down

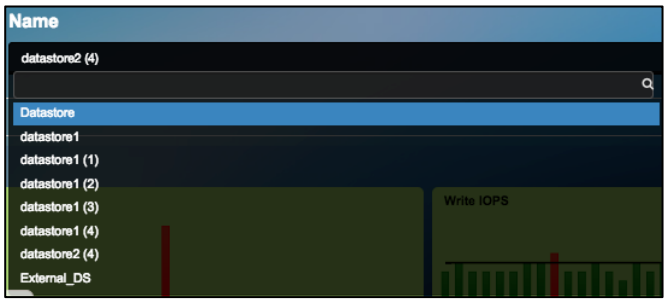

Figure 11.3: Names drop-down

Here is the Example of the Metric summary selected for VM 'Oracle_11g-n1' between 5:05am to 5:52am, when applications *postgres* and *mysql* performance are degraded, and where the root cause is pinpointed.

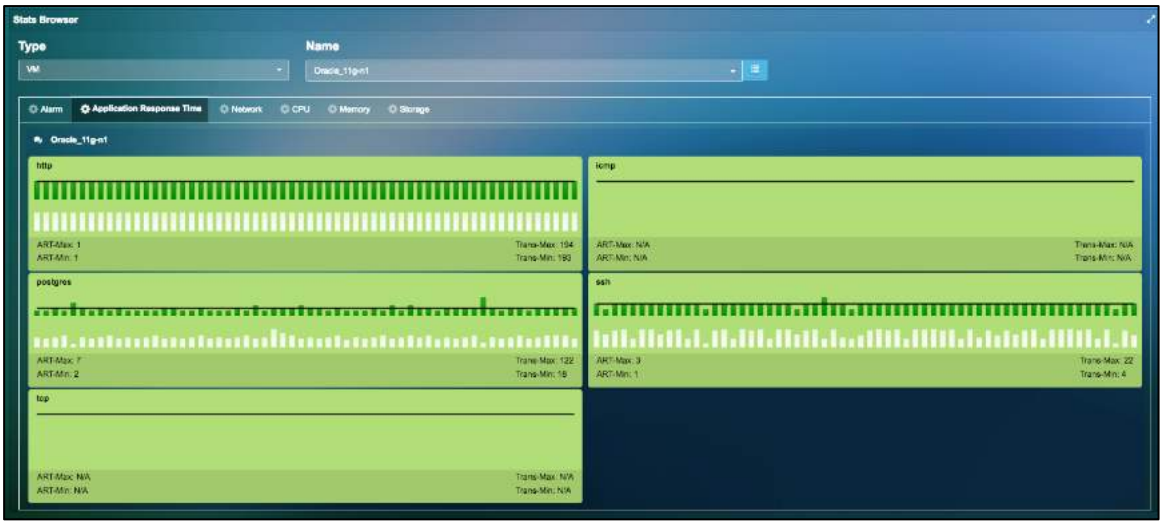

Figure 11.4: Application Response metrics for selected VM

Figure 11.5: Application Response metrics for selected VM

13. Alarms View

The alarms view provides a consolidated list of all alarms and provides their application response time (ART), uptime and their host in a tabular format. The alarms are sorted by severity and categorized into the five different categories application, network, CPU, memory and storage.

Figure 12.1: Alarms View

Clicking on the alarm will directly take us to the root cause view. Please see Section 13 for more information.

14. Reports

In order to view reports, click on the “reports” button the menu bar.

13.1. Report types

Uila allows you to either generate On-Demand reports or Schedule reports. You also have the option to generate reports in the CSV format.

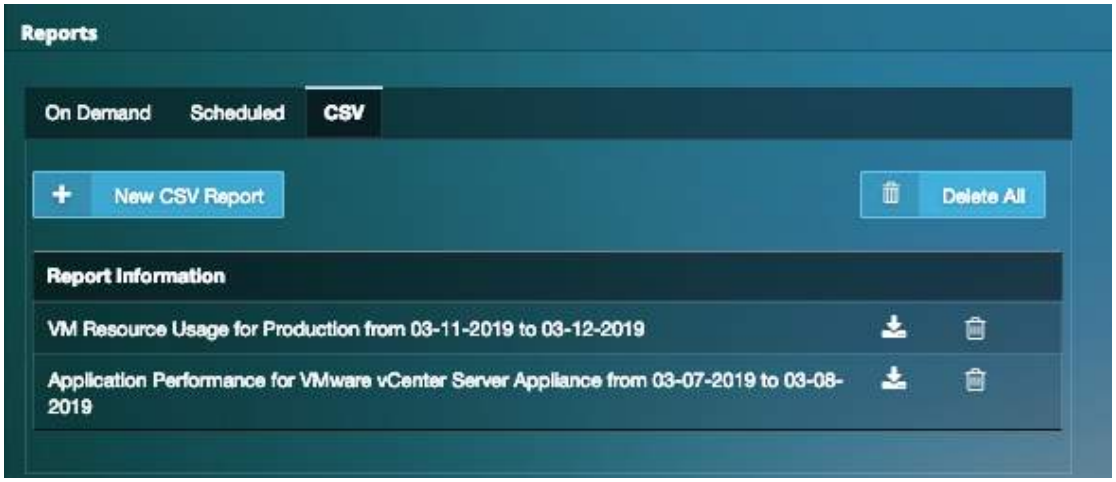

13.2. Report types

Uila has 4 different kind of reports

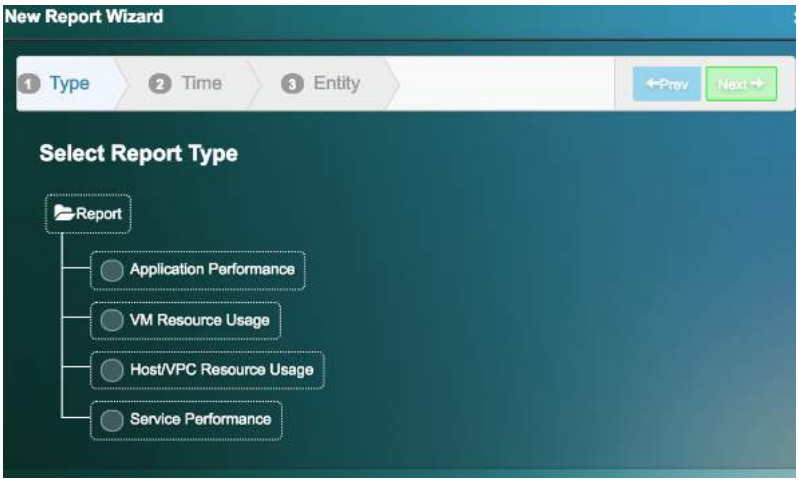

- **Application Performance** – Provides trend chats of the overall application performance of the entity selected (Datacenter, Cluster, Hosts or VM's) along with the CPU, Memory, Storage and Network.

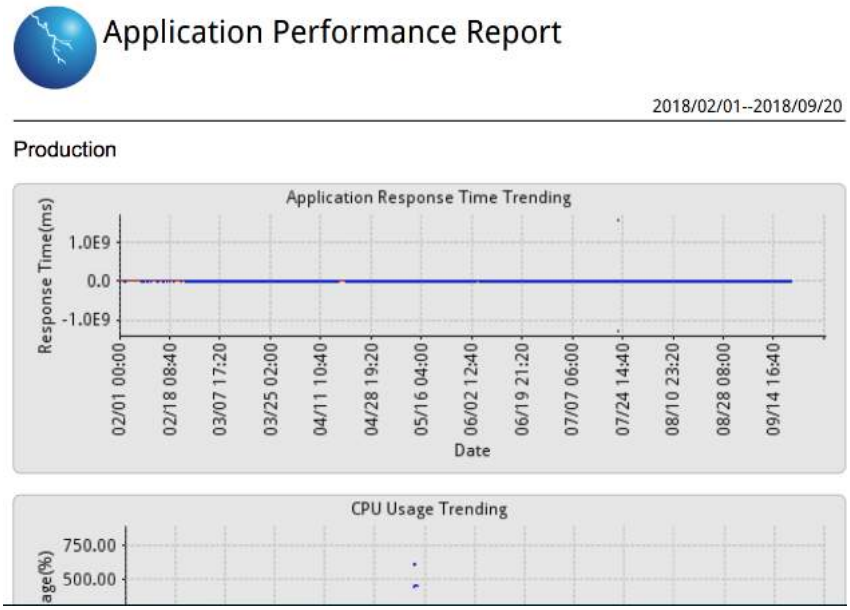

- **VM Resource Usage report** – With VM Resource usage report you can optimize cloud costs and coordinate between cloud governance teams and resource owners (IT teams) based on actual usage and uncover inefficiencies to reduce waste. You can visualize under-provisioned hosts or instances leading to application performance issue.

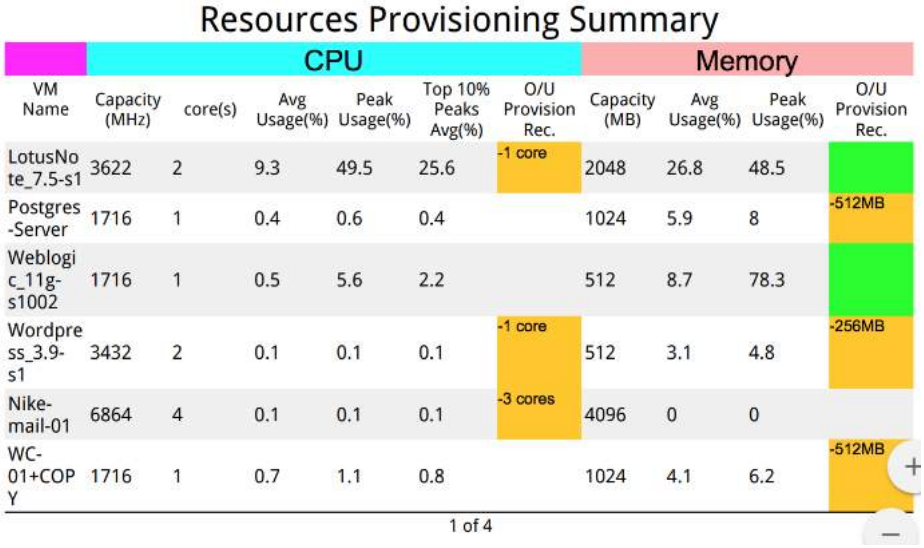

Resources Provisioning Summary

	CPU						Memory			
VM Name	Capacity (MHz)	core(s)	Avg Usage(%)	Peak Usage(%)	Top 10% Peaks Avg(%)	O/U Provision Rec.	Capacity (MB)	Avg Usage(%)	Peak Usage(%)	O/U Provision Rec.
LotusNote_7.5-s1	3622	2	9.3	49.5	25.6	-1 core	2048	26.8	48.5	
Postgres-Server	1716	1	0.4	0.6	0.4		1024	5.9	8	-512MB
Weblogic_11g-s1002	1716	1	0.5	5.6	2.2		512	8.7	78.3	
Wordpress_3.9-s1	3432	2	0.1	0.1	0.1	-1 core	512	3.1	4.8	-256MB
Nike-mail-01	6864	4	0.1	0.1	0.1	-3 cores	4096	0	0	
WC-01+COPY	1716	1	0.7	1.1	0.8		1024	4.1	6.2	-512MB

1 of 4

Please refer to the table below to understand the different colors in the Resources Provisioning Summary:

Resource (Color)	Provisioning	Peak Usage(%)	Top 10% Peak Ave(%)	Average Usage(%)
CPU (Orange)	OVER		< 50%	< 20%
CPU (Green)				20% ~ 60%

Resource (Color)	Provisioning	Peak Usage(%)	Top 10% Peak Ave(%)	Average Usage(%)
CPU (Yellow)				60% ~ 70%
CPU (Red)	UNDER			> 70%
Memory (Orange)	OVER	< 40%		< 30%
Memory (Green)		>= 40%		< 30%, or 30% ~ 80%
Memory (Yellow)		80% ~ 90%		
Memory (Red)	UNDER	> 90%		

- **Host Resource Usage report** – The host resource usage report provides the health summary of each hosts on its CPU, Memory, Storage and Network.

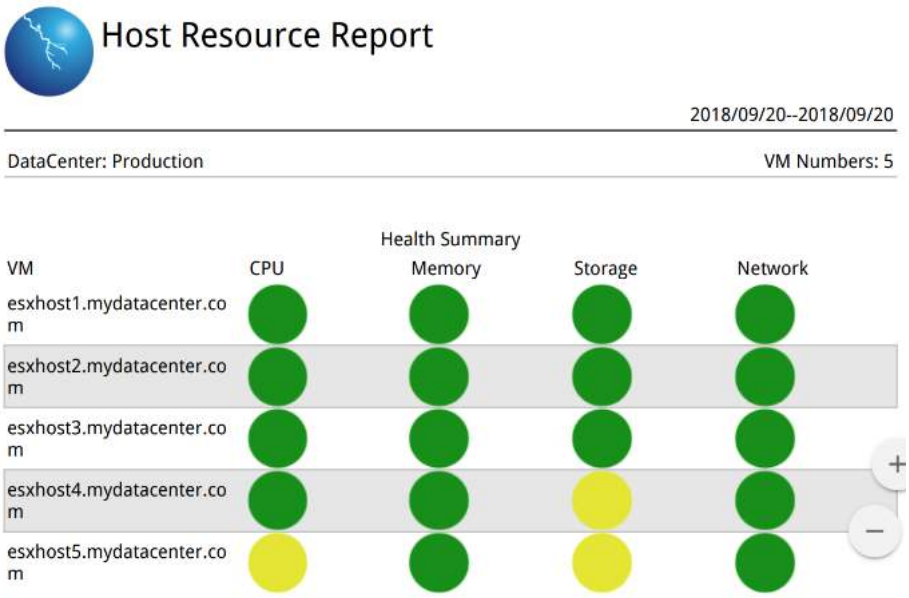

- **Service Performance Report** – The service performance reports provides the health of individual services running within the virtual machines.

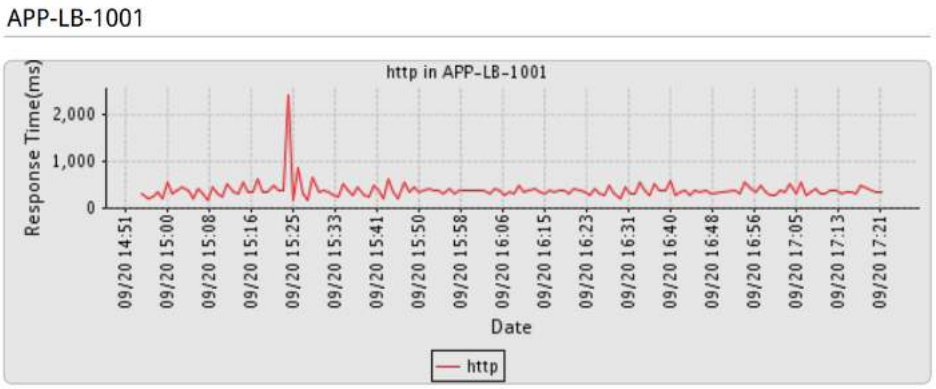

15. Appendices

15.1. Infrastructure and Application Statistical Counter for Measuring Key Performance Indicators

This Table summarizes all the statistical counters that Uila measured and collected from VMware vCenter or Hyper management server, and network packets, and stored in UMAS Big Data database:

Category	Counter	Type	Measurement Method	*Uila Built-in Best Practice Threshold (that overrides baseline value)
Application Performance	Application Response Time (ART)	KPI used for categorizing health score	Time (mSec) measured from the arrival of a client application request to the transmission of a server response.	Minimum ART baseline is 200 mSec. This means applications with less than 200 mSec response time will have Normal (green) ART health score.
	Network Round Trip Time (NRT)	KPI used for categorizing health score	Network round trip time (mSec) spent in the network	Minimum NRT baseline is 50 mSec. This means device with less than 50 mSec NRT will have Normal (green) NRT health score.

Network Infrastructure	TCP Fatal Retry	KPI used for categorizing health score	TCP re-transmit the same packet more than 3 times	No auto-learned baseline directly on TCP Fatal Retry packets. Health score is defined by the percent of TCP Fatal Retry count to total TCP packet count.
				If (x == 0) Normal
				If (0< x < 0.01%) Minor
				If (0.01% < x < 0.05%) Major
				If (x >0.05%) Critical

	Virtual Packet Drop (VPD)	KPI used for categorizing health score	# of Packet lost between vSwitch and virtual network driver	No auto-learned baseline directly on Virtual Packet Drops. Health score is defined by the percent of Virtual Packt Drops to total packet count.
				If (x < 0.01%) Normal
				If (0.01%< x < 0.05%) Minor
				If (0.05% < x < 0.1%) Major
				If (x >0.1%) Critical

	Zero Window	Statistics used for troubleshooting & investigation	TCP receive window closed. TCP receiver refused to receive more TCP data from the sender.	
	Reset	Statistics used for troubleshooting & investigation	TCP connection reset	
	Rx Bytes Average	Statistics used for troubleshooting & investigation	Number of bytes received	
	Tx Bytes Average	Statistics used for troubleshooting & investigation	Number of bytes transmitted	
	Usage Average	Statistics used for troubleshooting &	Number of bytes transmitted and received	

		investigation		
	Packets	Statistics used for troubleshooting & investigation	Number of network packets transmitted or received	
Storage Infrastructure	Disk Read Latency	KPI used for categorizing health score	Average mount of time (mSec) taken to process a disk read command	No auto-learned baseline for VM and Host Read Latency. Health score is determined by comparing to a fixed baseline value of 22 or 20 mSec for VM and host respectively.
	Disk Write Latency	KPI used for categorizing health score	Average amount of time (mSec) taken to process a disk write command	No auto-learned baseline for VM and Host Read Latency. Health score is determined by comparing to a fixed baseline value of 22 or 20 mSec for VM and host respectively.
	Kernel Latency	Statistics used for troubleshooting & investigation	Kernel average latency (KAVG) time an I/O request spent waiting inside the vSphere storage stack	
	Device Latency	Statistics used for troubleshooting & investigation	Device average latency (DAVG) coming from the physical hardware, HBA and storage device	
	Read I/O Ops	Statistics used for troubleshooting & investigation	# of Read operations per second	
	Write I/O Ops	Statistics used for troubleshooting & investigation	# of Write operations per second	

CPU Infrastructure	CPU Ready	KPI used for categorizing health score	Percentage (%) of time that the VM was ready, but could not get scheduled to run on the physical CPU due to physical CPU resource congestion	No auto-learned baseline for CPU Ready. Health score is determined by comparing CPU Ready value against fixed threshold below –
				For VM
				If (x < 5%) Normal
				If (5%< x < 10%) Minor
				If (10% < x < 20%) Major
				If (x >20%) Critical
				For host
				If (x < 10%) Normal
				If (10%< x < 15%) Minor
				If (15% < x < 25%) Major
				If (x >25%) Critical

	CPU Usage	KPI used for categorizing health score	Average CPU utilization (%) over all available virtual CPUs in the VM	No auto-learned baseline for CPU Usage. Health score is determined by comparing CPU Usage value against fixed threshold below –
				For VM
				If (x < 80%) Normal
				If (80%< x < 85%) Minor
				If (85% < x < 90%) Major
				If (x >90%) Critical
				For Host
				If (x < 85%) Normal
				If (85%< x < 90%) Minor
				If (90% < x < 95%) Major
				If (x >95%) Critical

	CPU MHz	Statistics used for troubleshooting & investigation	Average CPU MHz usage	

Memory Infrastructure	CPU Swap Wait Time	KPI used for categorizing health score	Average time (mSec) spent per minute a virtual machine is waiting for memory pages to be swapped in	No auto-learned baseline for CPU Swap Wait Time. Health score is determined by comparing CPU Swap Wait time percentage against fixed threshold below –
				For VM
				If (x < 300ms) Normal
				If (300ms< x < 1200ms) Minor
				If (1200ms< x < 3600ms) Major
				If (x >3600ms) Critical
				For Host
				If (x < 600ms) Normal
				If (600ms< x < 3000ms) Minor
				If (3000ms < x < 6000ms) Major
				If (x >6000ms) Critical

	Memory Active Usage GB/MB	Statistics used for troubleshooting & investigation	Amount of memory that is actively used, as estimated by VMkernel based on recently touched memory pages.	

	Memory Active Usage %	KPI used for categorizing health score	Amount of memory percentage that is actively used, as estimated by VMkernel based on recently touched memory pages.	No auto-learned baseline for Active Memory directly. Health score is determined by comparing Active Memory percentage (to total memory) against fixed threshold below –
				For VM
				If (x < 50%) Normal
				If (50%< x < 55%) Minor
				If (55% < x < 65%) Major
				If (x >65%) Critical
				For Host
				If (x < 40%) Normal
				If (40%< x < 45%) Minor
				If (45% < x < 55%) Major
				If (x >55%) Critical

	Memory Swap In Rate	Statistics used for troubleshooting &	Rate at which memory is swapped from disk into active memory	

		investigation		
	Memory Swapped	Statistics used for troubleshooting & investigation	Current amount of guest physical memory swapped out to the virtual machine's swap file by the VMkernel	
	Memory Consumed	Statistics used for troubleshooting & investigation	◦ VM: Amount of guest physical memory consumed by the virtual machine for guest memory. ◦ Host: Amount of machine memory used on the host. • Cluster: Amount of host machine memory used by all powered on virtual machines in the cluster.	

15.2. Reference Documents

VMware vCenter Installation and Setup Guide:
http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-552-installation-setup-guide.pdf

VMware Server and Host Management:
http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-552-host-management-guide.pdf

Virtual Machine Administration Guide:
http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-552-virtual-machine-admin-guide.pdf

VMware vSphere Monitoring and Performance:
http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-monitoring-performance-guide