



APPENSURE

## **Empowering Application Workloads Migration to Cloud Services**



APPENSURE

## Contents

|  |    |
|--|----|
| Introduction.....  | 2  |
| Cloud Migration Challenges .....   | 2  |
| Application Performance and Availability.....  | 4  |
| Application Workload Migration Challenges .....  | 4  |
| Incompletely Documented Applications .....   | 4  |
| Sprawling Application Architectures .....  | 4  |
| Architectures Not Designed For Elasticity.....   | 5  |
| Agile Development Process.....   | 5  |
| Essentials for Migrating Application Workloads.....                                      | 5  |
| Educated Application Placement.....  | 6  |
| Successful Migration Execution .....   | 7  |
| Continuous Application Performance Monitoring and Measuring in Migrated Environments.... | 7  |
| The AppEnsure Difference .....   | 8  |
| Summary of the key capabilities of AppEnsure .....                                       | 8  |
| Managing Cloud Migration with AppEnsure.....   | 8  |
| AppEnsure for Pre-Migration Discovery & Profiling of Workloads .....                     | 8  |
| Discover Applications.....   | 8  |
| Discover your Application Workloads Topology.....  | 9  |
| Measure Your Application Workload Performance.....                                       | 9  |
| Discover Every Tier of Every Multi-Tiered Application .....                              | 10 |
| Discover All Users of All Applications.....  | 10 |
| Discover Anomalies with Applications .....   | 10 |
| Profile Application Workloads.....   | 11 |
| Identify Work Loads to Migrate.....  | 11 |
| AppEnsure for Post-Migration – Continuous Monitoring for SLAs .....                      | 12 |
| Use Cases .....  | 12 |
| Large Media Conglomerate using AppEnsure to Profile Workloads on Cloud .....             | 12 |
| Medium Financial Firm Profiling Workloads to Migrate to Cloud.....                       | 13 |
| Conclusions .....  | 13 |




APPENSURE

## Introduction

With the increasing maturity and acceptance of cloud infrastructure services in the enterprise, migrating important workloads to the cloud—public, private or hybrid—is becoming a promising proposition. But migrating mission-critical applications, especially ones built on legacy architecture, can be challenging to say the least.

According to Citi Research, the market size of cloud computing, in terms of revenues, lies between \$45 to \$50 billion for 2013-2014, and will have a CAGR\* of 20% by 2020. Global Industry Analysts, Inc. (GIA), on the other hand, says the market for cloud computing will reach \$127 billion by 2017. As per the International Data Corporation (IDC), cloud software revenues will be \$67 billion this year, with a CAGR of 24% for the next five years. According to Gartner, Inc. (IT) enterprise spending on public cloud services will reach \$207 billion by 2016.

Public Cloud Services Market by Segment, 2010-2016



Source: Gartner (February 2013)

## Cloud Migration Challenges

The main concerns when thinking about cloud computing are:

**Security.** Moving your data and code to a third party provider creates some security risks.

**Loss of control.** For software-as-a-service (SaaS) and some platform-as-a-service (PaaS) solutions, the entire control of hardware, software, security policies, etc, is placed in the hands of a third party provider.

**Integration.** Most customers will need to integrate internal systems with cloud systems.

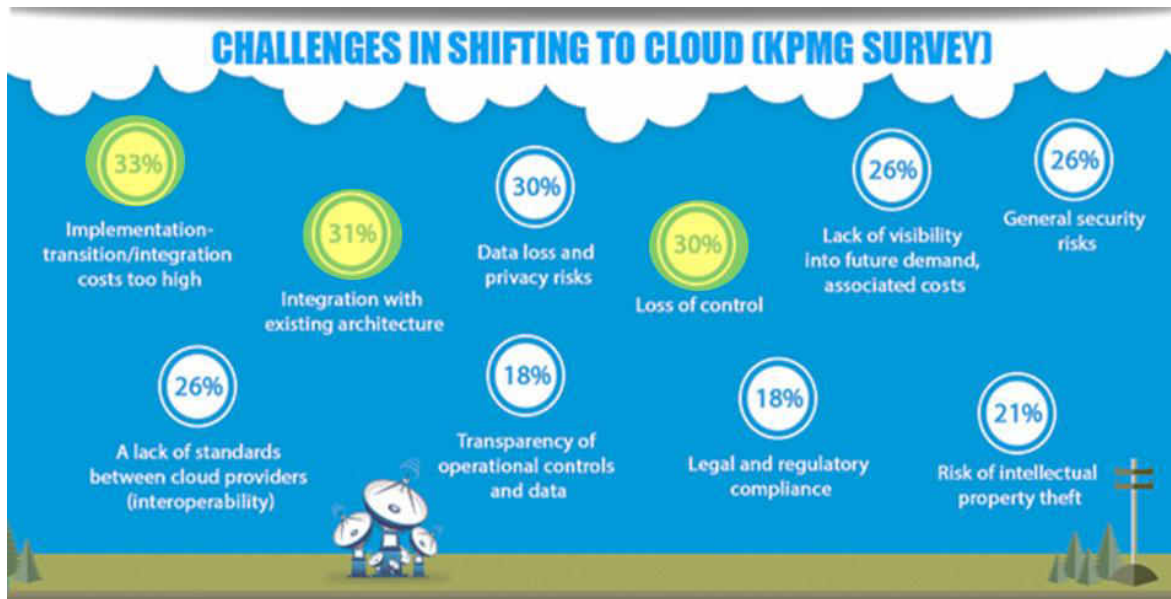
**Cloud service provider lock-in.** The concern is that once a cloud service of one provider is adopted, it will not be easy to switch to using an equivalent cloud service of a different provider.



APPENSURE


One of the top concerns for migrating seems to be the availability and control over the applications. Issues related to availability and reliability of cloud applications may arise from a combination of server performance, configuration errors, network design, and application architecture, possibly in combination, which can initially make them difficult to resolve.

A survey by KPMG, which recorded feedback from 674 respondents, reveals the following problems with adopting cloud computing:



Cloud Computing Trends: 2014 State of the Cloud Survey by RightScale involving 1,068 technical professionals across a broad cross-section of organizations outlines:

**Enterprises with Key Elements of Cloud Strategy Defined**



Source: RightScale 2014 State of the Cloud Report



## APPENSURE

## Application Performance and Availability

Existing enterprise applications must be thoroughly assessed to determine which workloads can benefit most from early migration to the cloud. Key considerations including costs of migration, **application redesign, application performance and availability**, security and privacy requirements, and regulatory requirements must be taken into account.

The application architecture will affect how an application can be migrated to public cloud. Enterprises will have to evaluate each application case-by-case to determine whether it's best-suited to a physical environment, virtualized environment, private cloud, public cloud or hybrid clouds. In every case, trade-offs will be required. Not all applications are architected to run in virtual or multitenant environments. In some cases, an application has intensive CPU or input/output needs that require a dedicated multithread processor and therefore don't lend themselves to virtual machines that share physical CPU and I/O capacity with other applications. In other cases, it may be because the application vendor has not structured the licensing of the application for a virtualized or multitenant environment.

Whatever the reason, IT decision-makers must first determine whether the application in question can run on a modern virtualized infrastructure "as is," whether it must be modified to do so, or whether the cloud is totally out of the question because of specific technical limitations.

The reason an existing application is still around is because the business still depends on it, and applications that the business depends on must deliver adequate levels of performance and availability.

Downtime of a critical application paralyzes a business. According to research by the Aberdeen Group, the estimated average cost of downtime is now \$138,888 per hour. The cost of downtime increased by 65% from 2010 to 2012. In 2005, according to Gartner, organizations lost \$42,000 every hour of application downtime and experienced 87 hours of downtime per year. In 2011, a Ponemon Institute study found that the downtime cost is \$3000,000 every hour and organizations had 14 experienced an average of 14 outages per year.



## Application Workload Migration Challenges

### Incompletely Documented Applications

Over time, developers often expand, maintain, and re-purpose applications, all without a shred of documentation. Worse yet, the applications rarely "retire" until a large-scale modernization or re-organization project such as virtualization or a private cloud initiative is undertaken. Lacking good information about the applications, migration teams often need to reverse-engineer and document the applications adding significantly to the project cost and time.

### Sprawling Application Architectures

Application craft is endemic. Acquisitions, fly-by-night cowboy coding and integration, all lead to spaghetti architectures and the proliferation of "silo" applications and dozens of incompatible platforms



## APPENSURE

and technologies – both legacy and new. Such architectures can dramatically increase the complexity of the migration effort. In many cases, companies may need to completely re-design the applications.

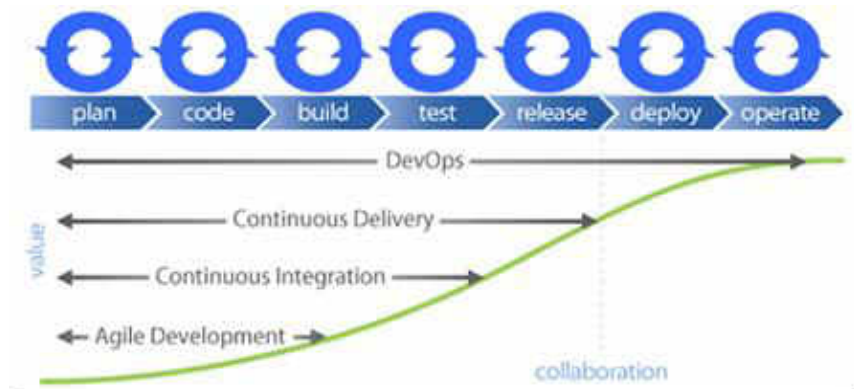
Even if existing single applications are based on standard multi-tier architectures, they still often have dependencies with other applications and shared data infrastructure that might or might not be compatible. For applications to work reliably in the target cloud environments, all application and data dependencies need to be considered. More complexity just means greater test effort to re-validate the systems in the cloud during the migration.

### Architectures Not Designed For Elasticity

Cloud architectures allow for elasticity: the ability to respond (scale out and in) rapidly to changing business demands. When the application load increases because of the number of users of the application increase, more resources (CPU, memory, storage, etc.) can be added dynamically to handle the load. Even if the underlying private cloud infrastructure supports scaling, the applications themselves may have to be overhauled to take advantage of this, adding to the complexity and cost of the migration effort.

### Agile Development Process

Often organization have agile development in process either by design or otherwise. Agile software development is a group of software development methods based on iterative and incremental development, in which requirements and solutions evolve through collaboration between self-organizing, cross-functional teams. It promotes adaptive planning, evolutionary development and delivery, a time-boxed iterative approach, and encourages rapid and flexible response to change. It is a conceptual framework that promotes foreseen tight iterations throughout the development cycle. Irrespective of the methodologies defined for agile development, often times “chaotic code” development takes place due to “need-this-by-Monday” deadline pressures.



### Essentials for Migrating Application Workloads

There are three major essential elements for migrating application workloads to public or private cloud services:

- Educated Application Placement
- Successful Migration Execution
- Continuous Application Performance Monitoring and Measuring in Migrated Environments



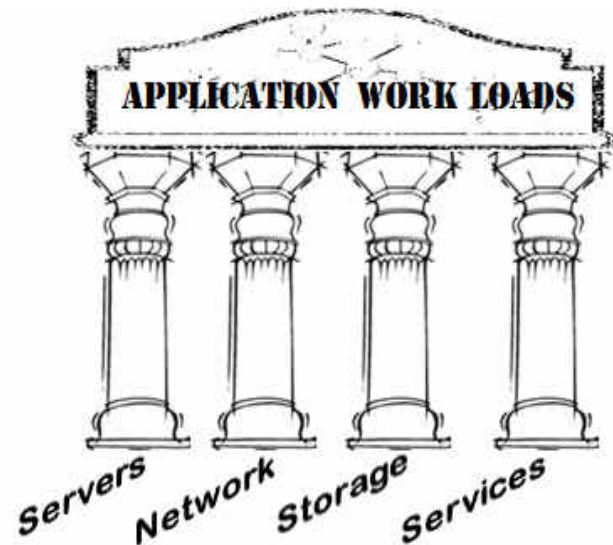
## APPENSURE

## Educated Application Placement

The first step in the cloud migration process is to discover current application workloads, the actual usage of application workloads to predict which (and how many) instance types are required. A disciplined real-time discovery process is required to truly understand an organization's application environment, including interdependencies, performance and utilization. Learning what applications are critical and type of performance is acceptable to end users can help determine where to put applications to migrate. Often times, there is significant risk and limited financial benefit of migrating legacy apps to the cloud.

Fundamentally applications are dependent on four resources: Server resources, Storage resources, Network resources and Services offered by infrastructure like DNS, LDAP, AD, etc.

For each application being migrated to cloud computing, consider the following application characteristics:



**Application availability.** The criticality of the application to business operations will determine the availability requirements that must be clearly specified in the cloud SLA.

**Application performance.** Depending on the performance requirements of the application, specific performance targets may need to be achievable with the cloud service.

**Application accessibility & security.** Moving an application to the cloud will require due diligence on the part of the cloud service customer to ensure proper security controls are in place and operating effectively.

Questions that need to be addressed for educated decision making are:

- What is the total application footprint in the organization and which of those are business critical and which of those are business line applications (like Mail Server, Office Suite, etc.)?
- What is the topology of each of those applications?
- What are the components that form each of the application?
- What are the dependencies (service, storage and compute power) of those applications?
- Who is using those applications?
- Why do they use them?
- How do they use them (to determine the application profiles)?
- When do they use it (to determine the load patterns)?
- What is the expected service level of these applications?
- What is the real service level being delivered today?

Organizations should understand the economics and suitability of the cloud service delivery model and how it fits specific application requirements. Using best practices like user sandboxes, proofs of concept



## APPENSURE

and bringing in network and security experts early in the decision process will also help ensure good choices and secure buy-in from stakeholders.

Application placement is an ongoing discipline that is based on understanding business, technical and workload constraints. The key is to always have the right workload in the right place at the right time.

### Successful Migration Execution

This step is often a process in and of itself. Most firms today offer complete migration services from simple lift and shift to full blown re-architecting of applications to take advantage of next generation cloud architectures. Once the correct applications have been identified and the cloud service delivery model determined, the migration and optimization of the applications can be planned and executed.

### Continuous Application Performance Monitoring and Measuring in Migrated Environments

Clouds obscure the operation of the infrastructure from the customer which makes it more difficult for the customer to ensure the performance of the applications running on those clouds.

Cloud Service Providers can differentiate themselves from Amazon by offering a tier of service that includes application performance service assurance.

It is very important to plan ahead of migration and implement monitoring and measurement solutions in the migrated environment that can independently (irrespective of the statistics the cloud provider offers) assess the performance and throughput of the workloads. Part of the challenge is the limited amount of service providers and management tools available today. It is critical to ensure the appropriate management and monitoring tools for the specific solutions chosen are in place and integrated to account for the existing infrastructure and the new cloud environment. It is essential to set up appropriate monitors, alerts and dashboards to ensure everything is performing up to expectations. Gaining deeper visibility into application workload performance allows you to create personalized read-only views for individual customers so they can always check the health of their deployment.

It is about more than just provisioning the right amount of compute, it requires monitoring and managing the end-user experience. Considerations must also be made about how to bring an application back to the internal environment if desired or if other external solutions become more compelling.

The process doesn't stop here. As demands on applications increase, as the cloud provider enhances their offerings, and as environments continue to evolve, it is critical to right-size deployments. You'll want to track and validate that the right instance types are in use, and the cost delta between on-demand and reserved instances. With that information in hand, deployments can be right-sized and the changes easily communicated back to the customer.

Making a cloud application placement decision and migration should always be part of a larger strategy around an organization's application and infrastructure portfolio. Cloud capabilities can provide significant benefits, especially if good decisions are made, migrations are planned and executed successfully, and early consideration is given to management of the hybrid environment.

AppEnsure offers solution that significantly help customers to ensure application performance in migrated environment by providing most meaningful and actionable intelligence about application workload topology, performance and behavior profiles so application placement is successful by benchmarking before and after migration.





APPENSURE

## The AppEnsure Difference

AppEnsure measures the response time and throughput for every application in your environment across the entire topology of each application system and uses unique analytics to provide meaningful and actionable root cause diagnostics. AppEnsure automatically identifies each application by name, provides a topology map for that application, and then provides a root cause analysis when response time, throughput, or the error rate for the application deviates from the norm. The AppEnsure solution works for all applications, irrespective of its source (custom developed or purchased) and irrespective of where it is located (physical, virtual, private cloud, or public cloud). With AppEnsure you can:

- Understand the reason for poor application performance for any and every application
- Stop wasting IT resources while in war room meetings guessing at the root cause of an incident
- Protect revenue and business productivity by preventing long outages

### Summary of the key capabilities of AppEnsure

**Instant Value to the Customer** – The ability to automatically discover, with absolutely zero manual configuration, all of the applications in the customer's environment, name them, map their topology across all of the tiers of the application system, and then in real time and on a continuous basis, measure their end-to-end response time and throughput. This includes the ability to measure the response time of the application from the perspective of the actual end user or from the perspective of the edge of the application system.

**Breadth of Applicability** – The ability to provide the value in #1 above for every Windows and Linux application that the customer owns, including both custom developed and purchased applications (90% of the applications that customers run are purchased). This includes the ability to follow an application as it migrates from physical hardware to a data center virtualization platform like VMware or Hyper-V, to a private cloud managed by a product like VMware vCloud Automation Center or OpenStack, to hybrid clouds hosted by vendors like Terramark or Savvis, and to public clouds like Amazon EC2, Windows Azure and various OpenStack based clouds.

**Automated Root Cause Analysis** – using a sophisticated combination of self-learning analytics and dependency rules, AppEnsure automatically tells the administrator what element in the infrastructure supporting the application is causing the degradation in the performance of the application or the downtime of the application. AppEnsure then proposes actionable resolutions that Ops can implement to return the application to the desired performance or required SLA.

## Managing Cloud Migration with AppEnsure

As mentioned earlier, educated application placement requires

- Discovering your applications and workloads
- Profiling your application workload landscape
- Defining SLAs based on application profiles
- Identifying the appropriate applications to migrate

## AppEnsure for Pre-Migration Discovery & Profiling of Workloads

### Discover Applications

AppEnsure uniquely and automatically discovers every Windows and Linux application in your environment, identifies it by name, and measures the end-to-end and hop-by-hop response time and



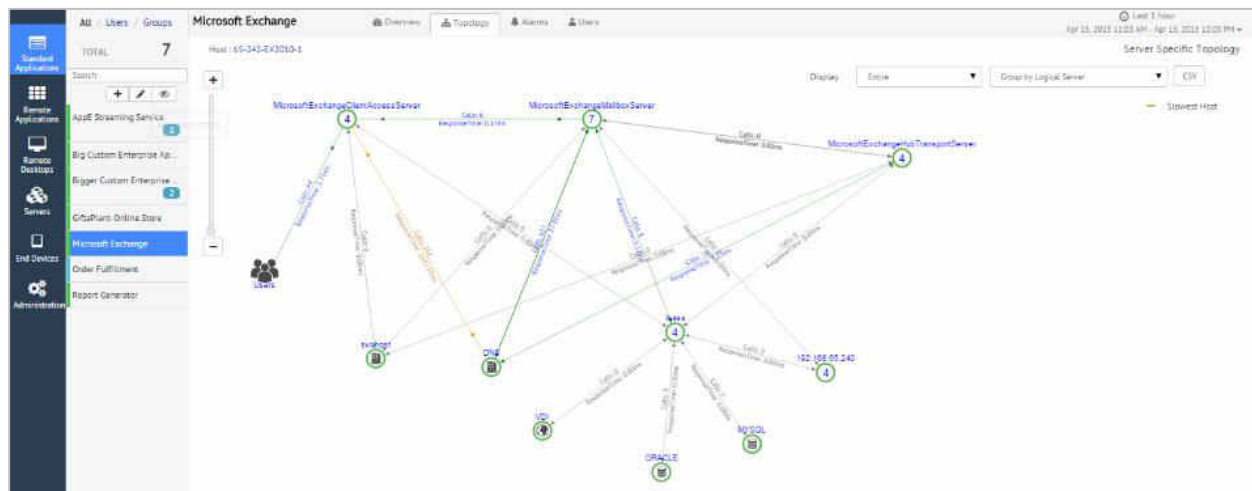
## APPENSURE

throughput of every application. This works for applications installed on physical servers, applications installed in virtualized guest operating systems, applications automatically provisioned in private or hybrid clouds, and applications running in public clouds. It also works irrespective of whether the application was custom developed or purchased.



### Discover your Application Workloads Topology

AppEnsure uses an auto discovery design, discovering application flows and mapping topology and user response times from these flows to create the application topology and update that topology in near real time as changes occur—all without user configuration steps. This allows for excellent “out-of-the-box” support for a large number of applications including custom apps. These capabilities significantly reduce the resources required to configure service models and operate the product.



### Measure Your Application Workload Performance

AppEnsure delivers an Enterprise view of all Apps running; legacy, custom & purchased in all locations; physical, virtual, private & public cloud. This is a dynamic view that will update in real time to show all instances with their response times. The overall throughput of all instances of an App delivers a deterministic demand load profile, not an inferred one from resource utilization. Measure and manage your performance!

AppEnsure provides visibility for each instance of Apps running by server with the specific performance data for that instance. The data allows the combination of App performance and server performance for



## APPENSURE

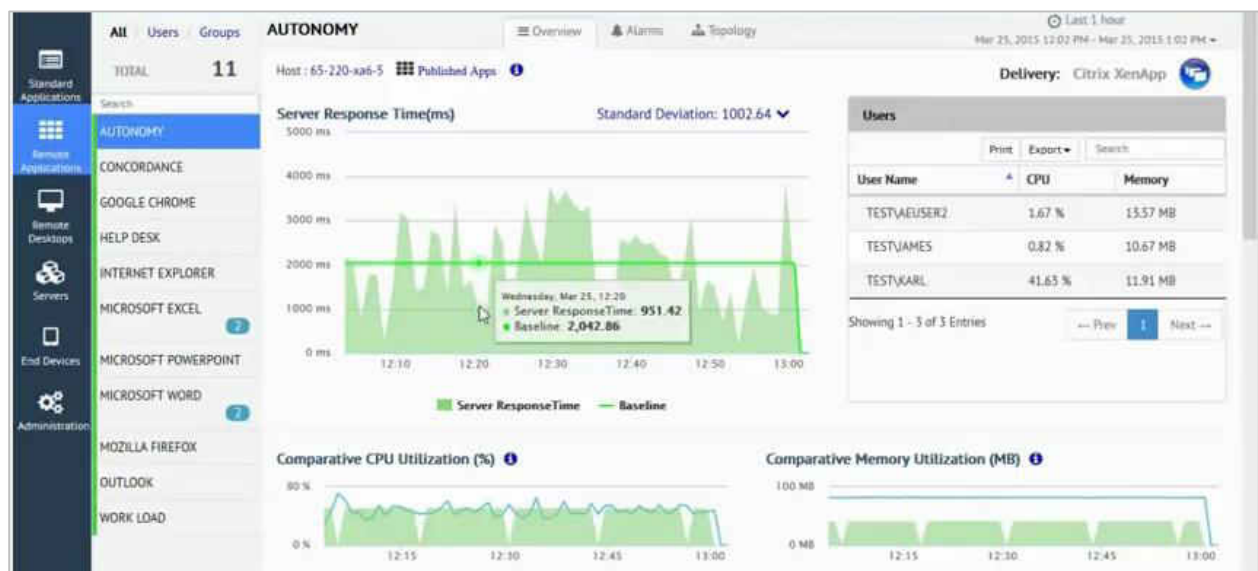
overall instance performance. As alarms are flagged, AppEnsure will in real time display the root cause and proposed remedy to return the instance to the desired performance or SLA.

### Discover Every Tier of Every Multi-Tiered Application

AppEnsure uses an auto discovers the different tiers for every multi-tiered application and provides the performance of each of the tier. This allows for excellent “out-of-the-box” support for a large number of applications including custom apps. These capabilities significantly reduce the resources required to configure service models and operate the product. AppEnsure automatically maps the virtual and physical topology of each App to give clear visibility of all transactions through all tiers of the App and infrastructure. Each tier is discovered, named and the transactional throughput and response times are shown. If there is an alarm for slow response time, the segment with extended latency will be clearly visible in the end to end transaction flow.

### Discover All Users of All Applications

AppEnsure identifies every user of every application and the experience that the user is having for every use of that application.



### Discover Anomalies with Applications

As deployed, AppEnsure uses a sophisticated anomaly detection algorithm to automatically assess when a response time excursion is valid, then deep diagnostics are triggered to analyze the event. In addition, the hop by hop segments latency is compared against the historical norms to identify deterministically which segment has extended latency and reduced App performance.

Based on many thousands of stored instances and rules, which are continually being incremented by machine learning, the root cause analysis of an event can be deterministically established. This triggers a proposed remedy for Ops to restore service to the SLA expected.



## APPENSURE

| Category      | Severity | Trigger       | Alarm Summary  | Time Frame                   |
|---------------|----------|---------------|--|------------------------------|
| Response Time | Minor    | Flow Analysis | MaxRequestPerChild is set to Zero. <span style="background-color: black; color: white; border-radius: 50%; padding: 2px;">118</span> | Wed Apr 15 19:08:21 2015 UTC |

| Diagnostic Detail   | Affected Systems | Time Stamp                   |
|---|------------------|------------------------------|
| Response time for Application: giftaplantonline is deviating because MaxRequestsPerChild is set to Zero, which may result in memory leakage. Swap memory : 2105 MB, No.of child processes : 1 | UAT-GAP-001      | Wed Apr 15 19:03:02 2015 UTC |
| Response time for Application: giftaplantonline is deviating because MaxRequestsPerChild is set to Zero, which may result in memory leakage. Swap memory : 2103 MB, No.of child processes : 1 | UAT-GAP-001      | Wed Apr 15 18:59:15 2015 UTC |
| Response time for Application: giftaplantonline is deviating because MaxRequestsPerChild is set to Zero, which may result in memory leakage. Swap memory : 2103 MB, No.of child processes : 1 | UAT-GAP-001      | Wed Apr 15 18:46:58 2015 UTC |

### Profile Application Workloads

AppEnsure Application Workload behavior profiles can be constructed using different time frames to view the performance of the workloads and creating reports for that time periods. This gives a full view of:

- Who is using those applications?
- Why do they use them?
- How do they use them (to determine the application profiles)?
- When do they use it (to determine the load patterns)?
- What is the expected service level of these applications?
- What is the real service level being delivered today?

Users can also define static thresholds to the response times and throughputs that they expect from application workloads. When response time and thresholds deviate from the thresholds set, then AppEnsure performs diagnostics to determine the root cause as to why those thresholds are not being met.

### Identify Work Loads to Migrate

By deploying AppEnsure in the current environment, you are now equipped with intelligence to make an informed decision. The length of time that you want to run AppEnsure solution in your current environment to gather all required varies and is determined by business cycle. Typically a 30 day interval will be a good time window to gather most of the information with your application portfolio.

At the end of this period you would have:

- Discovered all your Application Workloads
- Discovered all Application Workloads' Topology
- Discovered every tier of every Application Workload
- Discovered the services on which Application Workloads are dependent upon (Service Dependency Mapping)
- Measured every Application Workload's Performance in terms of response time and throughput that is provided to the users of the application (Service Level)
- Discovered all Users of all Applications
- Discovered the usage patterns of the Applications
- Discovered anomalies with Application Workloads



## APPENSURE

With this information, you can now start identifying the workloads that can be either fully or partially migrated to Cloud Services by addressing the following:

- Will the workload run in the target cloud environment? For example, compatible infrastructure, middleware and operating system image.
- Will the target cloud environment satisfy performance, availability and other non-functional requirements (NFRs)?
- Will the target cloud environment comply with applicable security, privacy and regulatory requirements?

Cloud computing is not an all or nothing scenario. Some criteria you should have while identifying the workloads to migrate:

- Stand-alone applications that do not interface with many other applications may be good candidates. Interfacing applications in the cloud with applications that still run locally can be problematic.
- Migrating applications that are already virtualized on a local server may be easier than those still run on dedicated servers.
- Fragile applications that continue to have problems when run locally are not good candidates for migration to the cloud. The existing problems are likely to get worse in the new environment.
- Latency can be a problem when running applications in the cloud.
- If you are planning on benefiting from scalability in the cloud, make sure that your software application is designed and configured to take advantage of it.

### AppEnsure for Post-Migration – Continuous Monitoring for SLAs

Clouds obscure the operation of the infrastructure from the customer which makes it more difficult for the customer to ensure the performance of the applications running on those clouds.

Cloud Service Providers can differentiate themselves from Amazon by offering a tier of service that includes application performance service assurance by leveraging AppEnsure.

You should be in a position to constantly monitor the Service Levels that are being provided by the applications once they are moved to the Cloud. Depending upon the APIs provided by the Cloud Service provider must not be the only way to measure the performance. You should tool your workloads to monitor the performance from an Application Centric view.

As demands on applications increase, as the cloud provider enhances their offerings, and as environments continue to evolve, it is critical to right-size deployments. You'll want to track and validate that the right instance types are in use, and you are getting the expected performance from the workloads.

### Use Cases

#### Large Media Conglomerate using AppEnsure to Profile Workloads on Cloud

A large media conglomerate is developing a CDN platform on public cloud to distribute video content on demand to its subscribers. AppEnsure is helping the customer to:

- Identify the workflows
- Identify the Service dependencies of the workloads
- Identifying the latencies of the workload so that architecture can be fine-tuned to accomplish target service delivery response times.



APPENSURE

### Medium Financial Firm Profiling Workloads to Migrate to Cloud

A medium sized financial firm (2000 users and about 100 workloads) is employing AppEnsure to:

- Profile the workloads – especially that involve published applications that interact with SaaS offerings like Salesforce
- Identify the bottlenecks in the current deployment
- Build a “true picture” in real-time the application service levels
- Identify which workloads should be migrated to the cloud

### Conclusions

Like any modernization initiative, there is no “one-size-fits-all” model to migrate Application Workloads to Cloud Services. A detailed profiling of the application workloads in real-time must be undertaken. As existing CMDB and other data bases that have tendency to be out of date with application workload proliferation, will not provide accurate picture of workload distribution.

AppEnsure solution empowers users to understand the real picture about the Application Workloads and helps in pre-migration process to identify the appropriate workloads that can be candidates for migration.

Since Cloud Service Providers obscure the operation of the infrastructure from the customer which makes it more difficult for the customer to ensure the performance of the applications running on those clouds, once migrated. Post migration, AppEnsure empowers customers to monitor service assurance of the applications running these clouds and help ensure the performance and identify issues with the Cloud Service Provider.