



CORPORATE SERVICE OFFERINGS - GLISTEN SOFTWARE PVT LTD

[HTTP://WWW.GLISTENSOFT.COM](http://www.glitensoft.com)

LAST UPDATED: 28TH NOVEMBER, 2017

ABOUT US

Our Vision

- Automate and Manage customer Infrastructure so that they can go to market reliably and quickly

Team

Team of 5 enthusiastic and dynamic folks having approximately 30 years of aggregated experience came together and provide DevOps solutions for bringing success to its customers

Office

Currently operating from Pune, India

ORGANIZATIONAL FOCUS


- OpenSource DevOps solutions to give customers cutting edge for their infrastructure operations and thereby obtain success. Our pillars to provide success are:
 - competency of latest technology stack
 - wide diversity of technology experience
 - customer focused delivery approach
- Implementation services for Cloud, IOT and BigData
- Data analytics and sentiment analysis implementation

SERVICE FOCUS




IT Ops

- Infra automation
- Operational automation



DevOps

- CI Automation
- Configuration management automation



Development

- NodeJS
- Java
- Python



BUSINESS FOCUS

Cloud Infrastructure DevOps

- Design and configure infrastructure in cloud
- Automate deployment of services through Continuous Integration (CI) and Continuous Deployment (CD)
- Aggregate logs from different services and servers for easy troubleshooting
- Application configuration and troubleshooting
- Identity Management and SSO implementation

Cloud Migrations

- Design migration strategy for least application downtime
- Execute migrations in optimal way
- Performance tuning
- DB migration

Automation

- Development of automated deployment of services and configuring them
- Automated security tuning
- Development Continuous Integration implementation

BigData

- Design, configuration and administration of BigData services
- Data pipeline configuration
- Performance benchmarking of services

IoT

- Configure and administer IOT infrastructure in Cloud
- Performance optimize services
- Monitor and manage infrastructure for high uptime

Data Analytics Development

- Sentiment analysis in Python
- Analytics on Social data collected from FaceBook, Twitter and LinkedIn to find social trends and behavior
- Analyze marketing attribution through data collection from Social sites

TECHNOLOGY COVERAGE

Virtualization / Cloud

- AWS
- MS Azure
- IBM SoftLayer
- OpenStack
- Docker
- Kubernetes

Configuration Management Development Tools

- SaltStack
- Ansible
- Ansible Tower
- Puppet
- Canonical - Juju
- Jira
- GitHub
- Jenkins
- CircleCI

Monitoring / Log Management / MessageBus

- Sensu
- Nagios
- ELK stack
- Kafka Cluster
- RabbitMQ Cluster
- Redis
- ActiveMQ
- Canonical - landscape

Data Analytics Development

- Python
- Java
- NodeJS

BigData

- MapR Hadoop
- Apache Hadoop
- Apache Drill
- Hbase
- Hive
- Spark / Spark streaming

SQL/NoSQL/TimeSeries DB

- MySQL
- PostgreSQL
- MS-SQL
- MongoDB
- InfluxDB
- OpenTSDB
- Aerospike

Others

- MAAS
- IPMI/SNMP management
- IAM
- Zookeeper Cluster
- Aerospike

ADDITIONAL IMPLEMENTATION EXPERIENCE

- Cisco Quad Implementation and integration with Cisco Universal Communication Manager (CUCM), Jabber and Cisco WebEx
- Implementation OpenStack services like Keystone, Glance, Nova, Neutron, Heat, Ceilometer and Swift.
- Handled openstack upgrade from IceHouse to Juno release
- Shibboleth SSO implementation
- Docker and Ansible integration on Raspberry Pi 2 for custom services

CLIENTS



CUSTOMER CASE-STUDIES

USE CASE – OPENSTACK PRIVATE CLOUD

About Customer

- Client has an ad-exchange platform that facilitates buying and selling of media advertising inventory from multiple ad networks
- With bidding happening in real time, client wanted to give low network latency services to its customers who were currently using AWS.

(IT-Dev) Ops need

- Maintaining Datacenter with 74 physical servers
- Automating the Infrastructure provisioning
- Automate OpenStack services configuration for quick horizontal scaling
- Key Challenges to be addressed:
 - Within few hours, new physical hosts can be added to production with minimal manual intervention
 - Complex add-on services were required to be deployed and administered

Implementation Details

- Hardware management and administration which included remote OS install, KVM console access, power-supply management using PDU. Managing 74 physical servers.
- Wrote SaltStack configuration management modules for OS configuration tuning, Installation of Applications like OpenStack, MapR Hadoop, Kafka, Zookeeper, Sensu monitoring tools.
- Administration of OpenStack services like Keystone, Glance, Nova, Nova-Network, Neutron, Heat And Ceilometer through SaltStack modules for easy horizontal scaling
- Administration of MapR Hadoop cluster for PAM-MySQL authentication, upgrades, troubleshooting. And services like Spark, Spark streaming, Drill, Hbase
- Administration of Kafka cluster and tuning. MirrorMaker configuration for replication across Kafka Clusters across regions
- Monitoring and Log-aggregation implementation for easy troubleshooting and pro-active reporting of incidents

USE CASE – AD EXCHANGE PLATFORM ITOPS/DEVOPS USING DOCKERIZED CONTAINERS

About Customer

- Client has an ad-exchange platform that facilitates buying and selling of media advertising inventory from multiple ad networks
- With bidding happening in real time, lot of distributed data processing and aggregation is required to analyze

(IT-Dev) Ops need

- Maintaining Datacenter with 74 physical servers
- Automating the Infrastructure provisioning
- Automating and maintaining the Big Data solution stack
- Key Challenges to be addressed:
 - 23 – 25 GB of data generated every day/210 TB in 2 months
 - 200 micro-seconds latency benchmark

Implementation Details

- Creating custom Docker images containing application services
- Writing docker files to build custom docker images.
- Configured Docker registry service to store custom images
- Custom bridge along with Docker network to configure static IP address for production usage. Bridge network allowed multi-host communication across docker containers.
- Docker containers talked to kafka cluster running on different host machines.
- Automated container configuration management using saltstack modules.
- Tested Docker orchestration using Docker Swarm and Apache Mesos

USE CASE – IOT APPLICATION CI/CD IMPLEMENTATION AND ADMINISTRATION

About Customer

- Client had IoT-based application in Car Telematics. Application device collected Car health data and sent it to a Cloud-based service every few seconds for showing different travel and health information related to the customer's car.


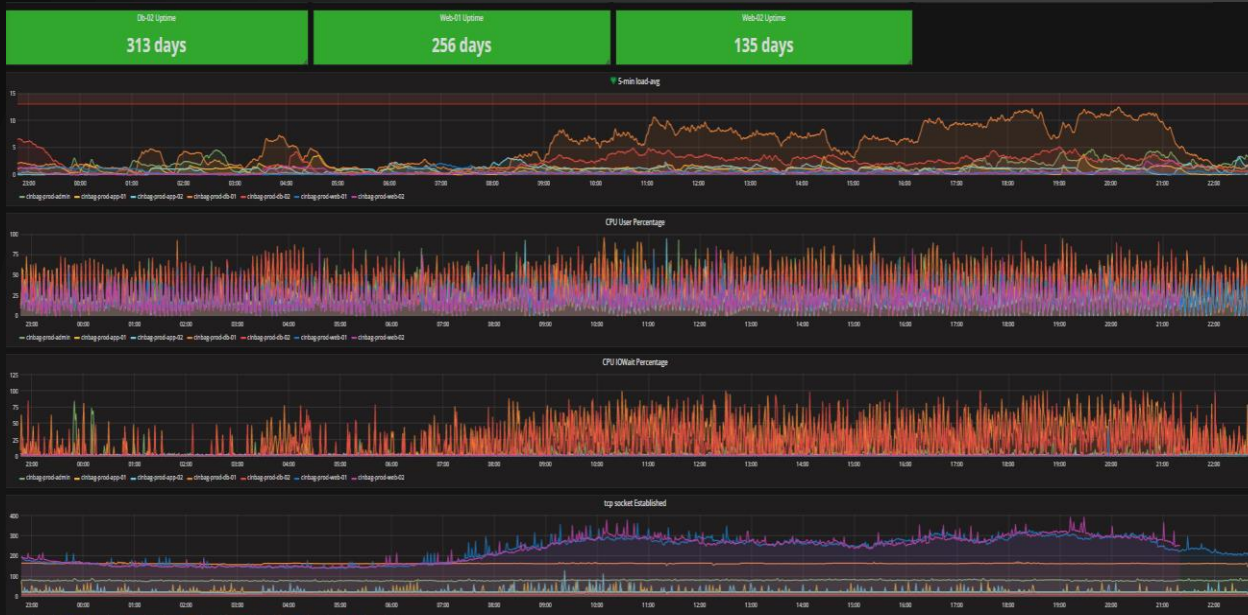
(IT-Dev) Ops need

- Different environments were distributed across MS Azure and IBM SoftLayer
- Deployment of complex and distributed services took days to onboard a customer
- Security, monitoring and log analysis was painful
- Migration from MySQL to MongoDB for geodata
- Performance benchmarking of the application was never done

Implementation Details

- To optimize quick deployment of environment, Ansible scripts were developed which configured all the application pre-requisites and configured new environment with single day
- Pro-active and advance monitoring and log-aggregation using ELK was implemented for easy troubleshooting for development team
- OpenLDAP and OpenVPN services were implemented to enhance security and access to the server. Security scan and auditing of the environment was done and issues resolved.
- Performance testing of the services was done to identify services having bottlenecks, so that horizontal scaling of application can be planned
- Custom Application API, Profiling and user pattern dashboards were created through ELK for application insight and usage
- Efficiently implemented MySQL and MongoDB clusters for storing application data

IOT APP MONITORING AND DASHBOARDING



CLIENTS > CINBAG-PROD-WEB-01



CINBAG-PROD-WEB-01
a few seconds ago
carlaq

_id carlaq/cinbag-prod-web-01
address 10.4.1.18
name cinbag-prod-web-01
redis {
 "host": "127.0.0.1"
}

silenced false
subscriptions check-cpu, check-disk, check-chrony, check-FD, check-load, check-memory-pct, check-web-ports, check-snmp, check-socket, check-alive, met-cpu, met-chrony, met-load, met-netif, met-socket, met-cpu-pct, met-memory, met-net, met-disk, met-netstat, met-mem-pct, met-uptime, met-pro-stat, met-interface, met-redis, client:cinbag-prod-web-01
timestamp 2017-07-27 22:52:40
version 0.26.3

Check

- check_memory_pct MEM OK - system memory usage: 44% a few seconds ago
- met_netstat cinbag-prod-web-01.tcp.UNKNOWN 0 1501176123 ... a minute ago
- check_web_ports CheckPort OK: All ports (6379,80,22,443,49999,8080) are accessible for host 0.0.0.0 a minute ago
- met_net cinbag-prod-web-01.net.eth0.tx_packets 1557751263 1501176177 ... a few seconds ago
- keepalive Keepalive sent from client 12 seconds ago a few seconds ago
- check_tomcat_running CheckPort OK: All ports (8080) are accessible for host 0.0.0.0 7 months ago
- met_uptime cinbag-prod-web-01.cpu.user 22169079.22 1501176170 a few seconds ago
- met_cpu_pct cinbag-prod-web-01.cpu.user 37.56 1501176146 ... a minute ago
- met_disk cinbag-prod-web-01.disk.sdc.reads 13069050 1501176176 ... a few seconds ago
- chrony-offset CheckChrony OK: NTP offset is 1.3501ms a few seconds ago
- met_netif cinbag-prod-web-01.eth0.rx_kB_per_sec 3.27 1501176123 ... a minute ago
- check_snmp CheckSNMP OK: All is well! a few seconds ago
- check_cpu CheckCPU TOTAL OK: total=41.62 user=37.56 nice=0.0 system=3.55 idle=58.38 iowait... 2 minutes ago
- check_chrony CheckChrony OK: NTP status is 'normal' a few seconds ago
- met_memory cinbag-prod-web-01.memory.total 7307198464 1501176175 ... a few seconds ago
- met_mem_pct cinbag-prod-web-01.memory.percent.free 3.3157248047067687 1501176175 ... a few seconds ago
- check_disk CheckDisk OK: All disk usage under 82% and inode usage under 85% a minute ago
- check_FD CheckFileR: 2336 allocated file handles a few seconds ago
- check_socket Usage: /opt/sensu/embedded/bin/check-socat.rb (options) -i,-input INPUT Input strea... a few seconds ago
- met_chrony cinbag-prod-web-01.chronystats.stratum 3 1501176179 ... a few seconds ago
- met_redis cinbag-prod-web-01.redis.uptime_in_seconds 22169042 1501176153 ... a minute ago
- met_socket network.sockets.total_used 672 1501176177 ... a few seconds ago
- check_load CheckLoad OK: Load average: 0.01, 0.28, 0.33 a few seconds ago
- chrony-stratum CheckChrony OK: NTP stratum is 3 a few seconds ago
- met_load cinbag-prod-web-01.load.avg.one 0.01 1501176177 ... a few seconds ago



USE CASE – CI/CD IMPLEMENTATION POC USING ANSIBLE TOWER AND JENKINS

About Customer

- Client had NodeJS application which they wanted to deploy automatically on Azure VMs in an automated fashion using Ansible Tower and Jenkins
- No manual intervention required

(IT-Dev) Ops need

- Application build should be automated
- Application deployment should be automated
- Success / Failure of deployment should be notified
- Integrate Ansible Tower with Jenkins

Implementation Details

- Configured Jenkins and AWX (opensource version of Ansible Tower)
- Integrated Ansible Tower with Jenkins, so that after building the application, an Ansible job is triggered for application deployment
- Ansible scripts are written which will read the build created by Jenkins and picked by Ansible Tower to be deployed on target host
- On successful completion of build and deployment, a success notification is sent
- On failure in building or deployment, triggers a notification with information details of job failure

USE CASE – MIGRATE APPLICATION TO DOCKERIZED CONTAINERS

About Customer

- Customer is creating search platform for various CMS platforms like Alfresco, Gmail, Exchange, Documentum, etc
- Customer is looking forward for easy horizontal scaling of environment

(IT-Dev) Ops need

- Horizontal scaling of application
- Configuring application services to function as micro-services
- Configure service configuration in distributed environment across multiple nodes
- Integrate with Continuous Integration to build docker images
- Clustering of MongoDB, Cassandra, ElasticSearch, Zookeeper, Kafka, and application services
- Monitoring of applications in distributed architecture

Implementation Details

- Wrote Dockerfiles to create different service images like MongoDB, ElasticSearch, Cassandra, Zookeeper, Kafka and application services
- Integrated in Continuous Integration to create Docker images and push it on Docker Hub
- Configured service restart for applications on start of container
- Wrote Ansible script to download images from Docker Hub, launch containers on different VMs, tune application configuration within Docker containers and restart services within containers
- Monitor application services
- Configure and tune service clustering for MongoDB, Cassandra, ElasticSearch, Kafka and Zookeeper

USE CASE – CONTINUOUS DELIVERY PIPELINE USING CIRCLECI AND AWS

About Customer

- Client is a book rental service provider in USA
- Since search and API service needs to be available always, there should not be any downtime.
- He has nodejs api services which need to be deployed to AWS in continuous delivery mode with 0 downtime.

(IT-Dev) Ops need

- Understanding of CircleCI build and deployment workflow
- Identify approach of app deployment for various conditions related to site failure and figure out solutions for 0 downtime
- Understanding of AWS services like Auto-scaling groups, load balancers and CI/CD pipeline

Implementation Details

- Configured CircleCI service to monitor production branch of GitHub. On changes committed to this branch, build is triggered.
- Build is stored on S3 bucket for ready processing
- Through custom python scripts called from CircleCI, provisioning is triggered in AWS
- Single instance of a load-balanced configuration is taken down for new code deployment.
- Post deployment, AWS services like load-balancers are configured and after new service node is functional then other node is taken down for maintenance. This ensures that application service is available 100%.
- MongoDB data is migrated to Elasticsearch for sub-second response time on search

USE CASE – AWS TO AZURE MIGRATION

About Customer

- Client had .NET application running in AWS Cloud with MS-SQL server as the backend
- Client wanted to migrate to MS Azure for more tight integration to Microsoft Cloud Technology Stack and optimize cost

(IT-Dev) Ops need

- With least downtime application should be migrated to Cloud
- Evaluate cost optimized solution available in MS Azure
- Easy deployment to MS Azure cloud
- Application security should not get compromised

Implementation Details

- Using Azure sites, customer was provided with a low cost and easy deployment solution ensuring secure access control
- MS-SQL DB migration was done through custom migration tools
- Dev, Staging and Production environments were kept in different resource groups, for secure access control
- Using Azure Slots, easy deployment and quick rollback capabilities were implemented for the application
- Auto-deployment of application through bitbucket was implemented through deployment tools available in Azure
- Utilized Mobile App and API App services of Azure for quick on-boarding of application

USE CASE – PERFORMANCE BENCHMARKING

- Write Ansible scripts for Apache Hadoop, Pig, Hive, benchmarking tools deployment
- Setup Apache Hadoop cluster and benchmark HDFS read/write operations using HiBench
- Benchmark Pig using PigMix and TPC-H
- Benchmark Hive using TPC-H
- Investigation on Pluggable Sort and Shuffle features of Hadoop 2.7
- Hbase benchmarking for insert, update and select queries

USE CASE – DATA ANALYTICS

- Analytics on social data, collected from FB, Twitter, LinkedIn to find social trend & behaviour to plan and improve marketing strategies
 - Functionality to view time series Sales trend report.
 - Functionality to hover on chart to see list of words that are causing the effect on Sales.
 - Functionality to Click on chart and see the summary of social content with sentiment score.
 - Functionality to update chart report based on date time search.
 - Functionality to compare two Sales charts of different date time.
 - Functionality to download report.
 - Functionality to send or share downloaded report.
- Collect Magento e-commerce data and apply analytics to get Marketing attribution
 - Functionality to view marketing attribute time series charts e.g. Clicks, Impressions, Sales.
 - Functionality to view chart of new and returning customer based on date and time filter.
 - Functionality to view chart of Impact on Sales for new vs returning customer.
 - Functionality to generate QoQ and YoY report along with sharing and downloading feature
- Development Stack
 - Python APIs – pandas, gensim, numpy, pymongo
 - ELK
 - Redis
 - MongoDB
 - AngularJS
 - Java
 - NodeJS
 - RabbitMQ
 - Glassfish server
 - SOLR

CUSTOMER RECOGNITION

- *Sagar and Glisten Software have done a tremendous job handling infrastructure for our HFC2 product, especially in areas of infrastructure automation, OpenStack Private Cloud and BigData. Sagar is a consummate professional and has been a real pleasure to work with on DevOps!*

- Tim McQuillen, Chief Knowledge Officer, Rubicon Project

- We have been working with Sagar and GlistenSoft for our connected car platform. The team brought a unique value by not only supporting in our complex developer ops requirements but also partnering with us to improve our solution through timely and helpful suggestions. We like to call this 'immersive partnership' what we have with GlistenSoft

- Vinu Kanakasabhpathy / Co-Founder and COO at CarIQ

THANK YOU

info@glistensoft.com