OpenCredo

MESOS

White paper:

Delivering Business Value with Apache Mesos

Executive Summary

In today's business environment, time to market is critical as we are more reliant on technology to meet customer needs. Traditional approaches to solving technology problems are failing as we develop and run applications at massive scale.

Apache Mesos is an open source project that addresses the problems of efficiently managing data center infrastructure at scale. Apache Mesos also provides a simplified abstraction layer for your development resources to use, allowing them to focus on delivering customer value.

In this white paper, we look at some of the core business drivers in the digital age, the problems of running and developing technology at scale and examples from disruptive companies like Ebay, Twitter and Airb&B who have Mesos in production.

The Problems of Technology at Scale

Development teams are spending more time dealing with increasingly complex infrastructures, rather than delivering on new product features. Computing resources in our data centers and the cloud are underused, increasing operating expenses. The demand for digital solutions at scale is growing by the day.

Introducing Apache Mesos

Apache Mesos efficiently manages computing resources at scale while providing a consistent model for developers to work with.

> *"WooRank reduced its infrastructure from more than 100 servers to 15. This has increased efficiency and resulted in savings of up to 45% on server costs."*

Benefits of Mesos

> *Dramatically reduce operating expenses by efficiently allocating workloads to make use of idle resources.*

Accelerate time to market as teams focus on customer needs, rather than configuration and infrastructure scaling.

Manage computing infrastructure at scale, without the need for complex configuration management processes.

Future proof technology investments. Apache Mesos supports key business requirements such as analytics and Big Data.

Case Studies

eBay increased developer productivity and reduced operational costs associated with virtualisation machine sprawl at eBay

eBay Inc - The worlds leading online auction site

eBay have heavily invested in continuous integration as a development best practice. At eBay, each developer gets a dedicated virtual machine running Jenkins, a widely used continuous integration server. This resulted in 1000's of VMs (Virtual Machines), which used physical resources and required operational support.

> *Continuous Integration is a development practice that requires developers to integrate code into a shared repository several times a day. Each check-in is then verified by an automated build, allowing teams to detect problems early.*

An internal analysis at eBay revealed that this massive VM footprint had fewer than 5% utilisation.

eBay developed a solution using Apache Mesos, a batch scheduling framework plugin for Jenkins and Docker, a container technology. The use of Docker, which can be used directly on Mesos with the Marathon scheduler, allowed eBay to support the complex tool chains and multiple environments, which their development teams use.

This use of Apache Mesos allowed eBay to remove the entire VM footprint from their infrastructure, drastically reducing operating complexity and costs, and freeing their PaaS team up to work on other areas.

Through clever use of Docker features such as the union file system, combined with Mesos and Jenkins eBay were also able to reduce build times dramatically - to a matter of minutes - increasing developer productivity.

Netflix developed a reactive stream processing service, which can collect and deal with millions of events per second, scaling up and down on demand

NetFlix is the world's leading internet television network

At Netflix operational insights are a core part of the business, and there is a relentless focus on observing these. The team at Netflix focuses on collecting data around areas such as system health, analytics and anomaly detection in real time to enable operational insights and enhance the overall customer experience.

OpenCredo

MESOS

White paper:

Delivering Business Value with Apache Mesos

Netflix has highly dynamic consumer demand, with workloads varying by orders of magnitude at different times in the day. These spikes and troughs have led to Netflix investing heavily in methods to dynamically scale up and down their infrastructure, reducing operational costs. Other systems, which act in unison with the core Netflix systems, must also scale dynamically.

By leveraging the plugin architecture of Apache Mesos, Netflix developed a custom framework and scheduler that meets their data stream processing needs. As a result, Netflix can now auto scale computing resources up and down based on the data processing needs at any particular point in the day.

OpenTable are accelerating developer productivity with MicroServices while reducing operational complexity and costs

OpenTable is a leading provider of hospitality solutions for restaurants.

Prior to 2012, OpenTable used a "big-bang" approach, with bi-weekly releases of a monolithic application. This approach was hampering OpenTables ability to innovate, and a decision was made to move to a development model supporting MicroServices. This gave developers much more freedom to use the right technology for specific tasks, but resulted in a massive increase in virtual machine usage. By mid-2013, OpenTable had over 1,200 virtual machines in use.

At this point, a large investment was made in configuration management using Puppet, which alleviated many of the problems, but still left OpenTable with a complex and expensive infrastructure to manage. Key development resources were being devoted to dealing with configuration management tasks rather than focusing on customer needs.

In early 2014, OpenTable started using Apache Mesos. Mesos allowed OpenTable to migrate many of their microservices from virtual machines by using the Marathon and Singularity frameworks. Apache Spark provides a Big Data platform and Apache Chronos allows OpenTable to run batch jobs when resources are under utilised on the Mesos cluster.

Developers have been freed from creating complex configuration management scripts and can focus on delivering customer value. As one OpenTable engineer pointed out at MesosCon;

> "configuration management is complex and hard, but if you have a good, simple way to package up and deploy your apps in a […] Docker Container it just makes everything so much simpler".

WooRank reduced infrastructure operating expenses by 45% in two weeks, whilst improving developer productivity.

WooRank analyses websites to provide clear, actionable insights on how they can be optimized.

Following their rapid growth since 2011, WooRank's infrastructure footprint had grown dramatically in both size and complexity.

This complexity was becoming a drain on WooRank's core technical team and their ability to quickly respond to customer needs.

WooRank was able to implement a Mesos based solution and migrate their existing infrastructure and processes in just two weeks. WooRank saw an immediate impact on both their ability to bring new features to market and operational efficiency. Apache Mesos allowed WooRank to consolidate their cloud infrastructure resulting in a 45% saving on operating expenses.

HubSpot increased developer productivity, rapid deployment and reduced operating expenses

HubSpot provides a cloud based integrated marketing platform

Hubspot's core focus is on improving developer productivity and freeing their development team up to focus on the problems they want to solve, rather than worrying about the infrastructure needed. At Hubspot, the development team of approximately 90 engineers deploys software up to 500 times per day, using a continuous integration and deployment model.

By 2012, Hubspot had over 2000 Amazon Web Services (AWS) instances in production, supporting production, staging and QA. By using Apache Mesos, HubSpot have been able to dramatically reduce the number of instances.

In their staging environment alone, HubSpot moved from 200 AWS instances of varying sizes to 18 large instances. Developers no longer need direct access to AWS as Apache Mesos takes care of requesting the correct resources to run their service.

The result: HubSpot have dramatically simplified their infrastructure, reducing operating expenses while increasing developer productivity.

Business Drivers in the Digital Age: Digital Innovation

Across multiple industries, trends such as mobile and the Internet of Things have already changed the way we approach business. We live in an era of heightened customer expectations. From smart cities to connected cars, location based services to the sharing economy, the need to deal with, and act upon, transactions, connections and data on a massive scale has never been greater.

The Internet of Things is creating massive business opportunities and huge scalability challenges. In a recent report, Goldman Sachs predicted 28 billion devices will be connected by 2020, a 10X increase from today.

OpenCredo

MESOS

White paper:

Delivering Business Value with Apache Mesos

Consumers and businesses now expect rapid and iterative evolutions in technology. The traditional "Big IT" approach to technology delivery, where timescales were considered in terms of years is no longer an option, as customers will move to a competitor long before a traditional technology project can get through the planning phase.

The most valuable digital resource in your organisation is your technical talent. Too often they can be tied up with operational tasks, rather than tasks with a customer value. A technology to ease this operational burden, allowing your technologists to focus on creating customer value is now a business necessity.

Operational Efficiency

The new breed of digital applications, running at massive scale and supporting customers run into the millions and need managing in a different way. Cloud computing and virtualisation have given businesses access to unparalleled levels of computing power. But with such access comes complexity and inefficiency in our operational approaches.

Simply put, businesses are leaving margin in the data center without ever realising it.

The promise of cloud computing and virtualisation has long been workload consolidation and increased efficiency. However, the established paradigm remains, and virtualised instances of machines are viewed as, and treated like, individual servers by both operations and development staff.

We treat our cloud and virtualised servers as we did physical servers before. We create static partitions, segregate servers and workloads and focus our capacity planning around peak usage. We accept our resources being under utilised, and the additional operating expenses that this entails.

In 2008, McKinsey reported on very low levels of utilisation in data centers. More recent research in 2013 from Accenture demonstrated that the problem still remains in public cloud environments, with overall utilisation hovering around the 10% mark. Your infrastructure investment is not working as hard or as efficiently as it could.

Emerging development best practices such as MicroServices are already exacerbating virtual server sprawl as people continue with a one server per service deployment model. Even as we decrease capital expenditure by moving to the cloud, we continue to see our operating expenses and complexity increasing.

MicroServices are an approach to developing a single application as a suite of small, connected, services.

Introducing Mesos

What if a technology existed that allowed you to manage your infrastructure at scale, reducing complexity, operating expense and allowed your most valuable technology resources, your people, to focus on delivering the best customer experience possible?

Apache Mesos is an open source technology that allows you to manage your infrastructure at scale, while reducing complexity and operating expenses. In turn, allowing your most valuable technical resource, your people, to focus on delivering the best customer experience possible.

Disruptive companies that are redefining entire industries, such as eBay, Twitter, Netflix, OpenTable, HubSpot, WooRank and AirBnB already have Mesos in production.

Mesos gives you the flexibility to re-assign under used resources as and when you need too. Mesos allows you to view and manage your data center in its entirety. You can manage your usage of technology and match this usage to your business needs, from dealing with capacity issues at peak times to running Big Data technologies such as Hadoop overnight.

A large number of frameworks have already been created for Apache Mesos, allowing developers to rapidly deploy Big Data, Batch Scheduling, Data Storage and PaaS tools/long running services across the same cluster of machines.

McKinsey noted developers are spending up to 30% of their time on systems engineering tasks rather than software development.

About OpenCredo

OpenCredo is an expert team of Software Development consultants. We specialise in Big Data, NoSQL and Cloud Computing.

info@opencredo.com | 5-11 Lavington St, London, SE1 0NZ | T. 020 3603 2680 | www.opencredo.com

References

O'Grady, Stephen (2013), The New King Makers, O'Reilly

Liu, Haun, 2013 A Measurement Survey of Server Utilization in Public Clouds, Accenture

Ebay (2014) Delivering EBays CI Solution with Apache Mesos, [Online], Available: http://www.ebaytechblog.com/2014/04/04/delivering-ebays-ci-solution-with-apache-mesos-part-i/

Fowler and Lewis (2014) Microservices, [Online], Available: http://martinfowler.com/articles/microservices.html, ThoughtWorks

Goldman Sachs (2014), The Internet of Things: Making sense of the next megatrend, Goldman Sachs, [Online], Available: http://www.goldmansachs.com/our-thinking/outlook/internet-of-things/iot-report.pdf

McKinsey (2008) Data Center Survey, McKinsey Insights

McKinsey (2013) IT Trends Survey for digital innovation, McKinsey Insights

McKinsey (2009) Where IT Infrastructure and business strategy meet, McKinsey Insights

Netflix (2014) Reactive Stream Processing Using Apache Mesos, [Online], Available: http://youtu.be/T1-cQ_F2zZo

EBay (2014) Delivering EBays CI Solution with Apache Mesos and Docker, [Online], Available: https://www.youtube.com/watch?v=VZPbLUJnR68

OpenTable (2014) Mesos @ OpenTable, [Online], Available: http://www.slideshare.net/samsalisbury/mesos-at-open-table-40675874

HubSpot (2014), Continous Delivery at a rate of 500 Deployments per day, [Online], Available: https://www.youtube.com/watch?v=bIjdZO1g5ZE

Thoughtworks (2014), Continuous Integration, [Online], Available: http://www.thoughtworks.com/continuous-integration

About OpenCredo

OpenCredo is an expert team of Software Development consultants. We specialise in Big Data, NoSQL and Cloud Computing.