Tech Talk - **Devops – a New Software Development Paradigm**

June 23, 2015

DevOps

Exploiting the natural tension between Development & Operations

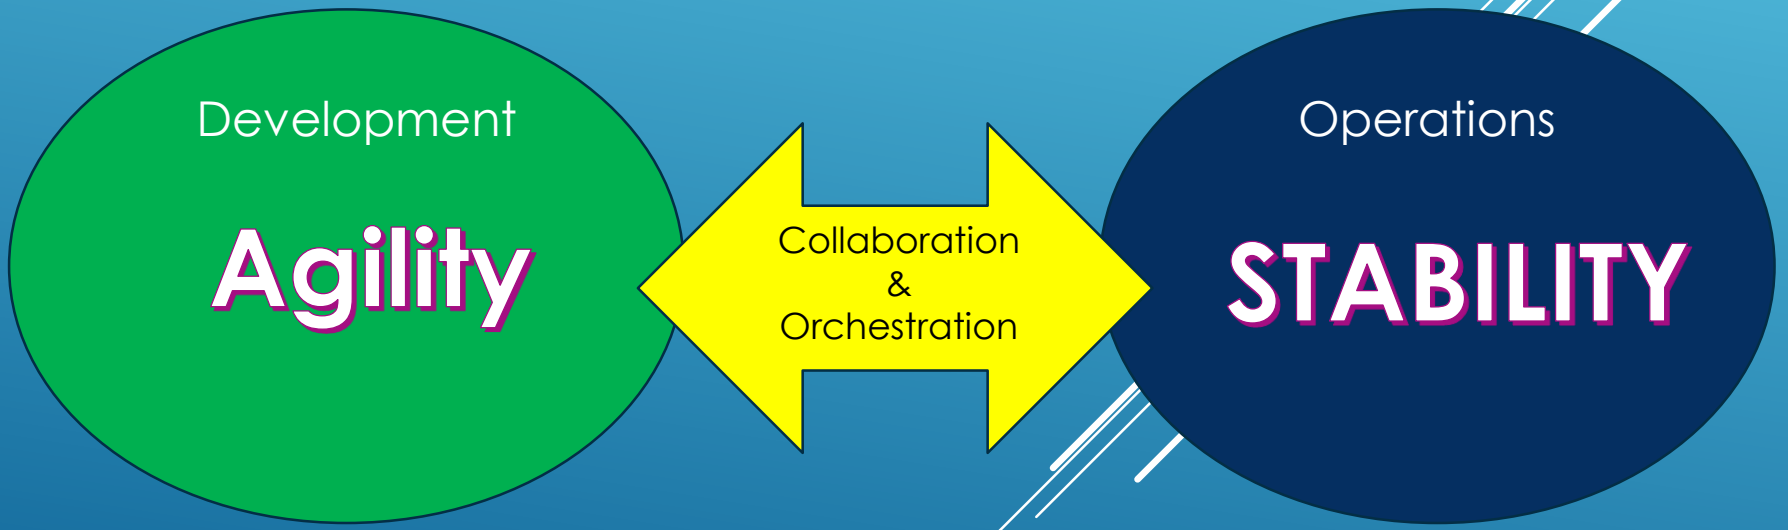

DevOps

a methodology of software development that emphasizes communication, collaboration and integration between software developers and operations personnel.

a response to the interdependence of software development and IT operations.

It aims to help an organization rapidly produce software products and services

DevOps History (per Wikipedia)

- At the Agile 2008 conference, Andrew Clay Shafer and Patrick Debois discussed "Agile Infrastructure", afterwards creating the Agile System Administrators Group on Google.

- The term "DevOps" was popularized through a series of "DevOps Days" starting in 2009 in Belgium. Since then, there have been DevOps Days conferences held in many countries worldwide.

A DevOps Functional Definition

Agile Methods applied to Software Engineering

DevOps requires a Lean practice

What is Lean?

- Maximizing customer value while minimizing waste.
- The goal is to provide perfect value to the customer through a perfect value creation process that has zero waste.

- If an effort doesn't add value, don't do it.

DevOps Simplified Phases

- Release
- Configuration Management
- Orchestration
  - Continuous Integration/Continuous Delivery
- Monitoring

DevOps vs traditional Agile

Devops = Agile + Process Rhythm

DevOps Getting Started

- Tools Inventory
- Skills Assessment
- Task to be accomplished
- Tool Selection (there are many more tools than can be used!)
- Assignment of Responsibility

- DevOps should be language neutral, although modern languages popular today are: Java, Ruby, Python, C++,…

DevOps lifecycle – (New Relic)

- Create new code changes
- Check in code
- Pull code changes, check latest patches, and build
- Test (continuous integration server to generate builds & arrange releases):
  - individual models
  - integration tests
  - user acceptance tests.
- Store artifacts and update the build repository (repository for storing artifacts, results & releases)
- Deploy and release (release automation product to deploy apps)
  - Configure environment
  - Update databases
  - Update apps
  - Push to users
- Application & Network Performance Monitoring (preventive safeguard)
- Repeat

DevOps Tools & examples

- release (jenkins, travis, teamcity)
- configuration management (puppet, chef, ansible, cfengine)
- orchestration (zookeeper, noah, mesos)
- monitoring, virtualization and containerization (AWS, OpenStack, vagrant, docker)

DevOps course materials – written with Chancellor Pascale

1) OS, Virtualization, & IAAS
2) Software development
3) Testing
4) Configuration Management
5) Deployment
6) Monitoring and Alerting
7) Caching and Queues
8) Web and Application Servers
9) Security
10) Dashboards

Module 1: OS, Virtualization

OS's
- Linux
- Windows

Virtualization
- XenServer, VirualBox, KVM, ESX/ESXi, Hyper-V, Docker

IAAS
- Google Compute Engine
- Amazon EC2
- Microsoft Azure
- RackSpace
- JoyentCloud

Linux in the Cloud

- Most common cloud-deployed Linux distros
  - Red Hat Entreprise Linux
  - SUSE Linux
  - Mint
  - Ubuntu
  - CentOS
- Support virtualization and can be easily virtualized
- Can be targeted as client or server with variable feature sets and footprints

Virtualization - XenServer

- Core of this product is Xen Project, which is an open source project that provides a bare metal virtualization
- XenCenter product provides ability to manage virtual machines across numerous physical machines
- Role-based administration allows for access controls on clusters of VMs with hierarchical access across hypervisors.
- Management console runs on Windows desktop OS

Virtualization – VirtualBox

- Another open source software option managed by Oracle
- Can virtualize almost any modern operating system
- Runs on Windows, Linux, Mac OS X, etc.
- HyperBox is an enterprise virtualization manager that utilizes VirtualBox at its core

Virtualization – KVM

- Linux-based virtualization solution running on x86 hardware
- Run Linux or Windows virtual machines
- VMs have personal virtualized hardware (network card, hard disk, graphic card, etc.)
- Requires QEMU

Virtualization – ESX/ESXi

- vSphere ESXi Hypervisor is a purpose-built bare metal
- Simple configuration makes deployment quick and easy
- vMotion provides a mechanism for transferring virtual machines between hypervisors on different machines
- Includes stateless firewall for hypervisor allows for setting firewall rules on individual IPs or IP ranges.

Virtualization – Hyper-V

- Microsoft-based hypervisor that runs on Windows Server 2012
- GUI and command line-based management tools
- Support virtual desktop infrastructure (VDI)

Virtualization – Docker

- Container provisioning and management
- Distribution mechanism for VMs and applications
- Build clusters to run on top of hypervisor
- Sits on top of virtualization management software like VirtualBox
- Designed to allow for simple provisioning of a large number of purpose built containers.

Module 2: Developing Software

- Developing Software
  - Codenvy
  - Koding
  - Cloud9
  - JSFiddle
  - PythonFiddle
- Continuous Integration
  - TravisCI
  - AppVeyor
  - CloudBees
  - CodeShip
  - Git/GitLab/GitLab-CI
- Deploy
  - Heroku
  - OpenShift
  - Google App Engine
  - AWS Elastic Beanstalk
  - Unity3d

CODENVY

- Integration with Docker via Dockerfiles
- Provision private cloud for your own enterprise
- IDE supports projects in PHP, Python, JS, Go, Java, Scala, C++, Ruby, etc.
- Running projects have url and can be run for free with low usage, great for start-ups and teams
- Incorporates Git support
- Easy to use tools for building and running instances

KODING

- Supports PHP, C++, Python, Ruby
- Direct integration with GitHub
- Access to VM that holds your code directly through full featured terminal application
- VMs can include multiple applications that work together.
- Simple interface similar to IntelliJ IDEA
- Hints of social media, with channels for communications amongst its users

CLOUD9

- Integrates with GitHub, allowing cloning of any GitHub project into its editor suite
- Exposing more control of VMs including sudo access, which allows for customized installation of services and packages.

JSFIDDLE

- Allows for editing of HTML, CSS, and JavaScript simultaneously and presents the page as it would be render in the browser.
- Built around sharing and collaboration via sharing of links to code and saving variants.
- Allows development in a plethora of JS frameworks from windowing to data visualization
- Can simulate external APIs in HTML, JSON, JSONP, and XML
- Hinting and code clean up options are available

PYTHONFIDDLE

- Similar to JSFiddle but for python
- Less options for external inclusion, only dozens of packages
- Collaboration and sharing via link sharing and inclusion of StackOverflow answers

TRAVISCI

- Web-based continuous integration platform
- Builds code from a variety of repositories, including direct integration with GitHub users and projects
- Works in most modern languages, such as Java, C++, Ruby, Python, etc.
- Create travis.yaml to configure the build including tests that reach out to a number of services and databases.
- Travis builds run in either Linux or Mac OS X virtual machines
- Dashboard shows any code updates that trigger builds, any comments for this revision, and the outcome of builds.

APPVEYOR

- Integrates with BitBucket and GitHub projects
- Build .Net projects with little to no configuration
- The ability to build, test, and deploy from within this one service
- Can deploy directly to Azure or Amazon AWS
- Package management and distribution through NuGet
- Ability to run code on bare metal machines with SSD storage.

CLOUDBEES

- Cloud-based Jenkins Continuous Integration
- Workflow service allows for development of pipelines for distribution of simple and complex projects
- Integrates with PAAS providers like GAE, AWS Elastic Beanstalk, and Pivotal Web Services
- Deploy to Compute Engine, AWS, and Verizon
- Test and deploy mobile applications
- Includes enterprise security features such as VPN connections and LDAP managed access

CODESHIP

- Integrates with BitBucket and GitHub
- ParallelCI service runs test suites in parallel to make builds faster.
- Deploy to services like Heroku, AWS, etc.
- Includes free community version for open source projects
- Multiple pipelines for quality assurance and production
- Ability to log onto debug builds via SSH
- Supports Ruby, Node, PHP, Python, Java and Go

GIT

a distributed revision control system with an emphasis on speed, data integrity, and support for distributed, non-linear workflows

WHAT IS THE DIFFERENCE BETWEEN GIT AND GITHUB?

- Git is a version control system; think of it as a series of snapshots (commits) of your code. You see a path of these snapshots, in which order they where created. You can make branches to experiment and come back to snapshots you took.

- GitHub, is a web-page on which you can publish your Git repositories and collaborate with other people.

GIT

http://www.git-scm.com/ — Git — Understanding the basics of Gi...

File Edit View Favorites Tools Help

git --everything-is-local

Search entire site...

Git is a free and open source distributed version control system designed to handle everything from small to very large projects with speed and efficiency.

Git is easy to learn and has a tiny footprint with lightning fast performance. It outclasses SCM tools like Subversion, CVS, Perforce, and ClearCase with features like cheap local branching, convenient staging areas, and multiple workflows.

Learn Git in your browser for free with Try Git.

About
The advantages of Git compared to other source control systems.

Documentation
Command reference pages, Pro Git book content, videos and other material.

Downloads
GUI clients and binary releases for all major platforms.

Community
Get involved! Bug reporting, mailing list, chat, development and more.

Pro Git by Scott Chacon and Ben Straub is available to read online for free.
Dead tree versions are available on Amazon.com.

Windows GUIs Tarballs
Mac Build Source Code

Companies & Projects Using Git

Google facebook Microsoft twitter LinkedIn NETFLIX PostgreSQL

android RAILS Qt GNOME eclipse

GITLAB / GITLAB-CI

- Open source Git-based repository and continuous integration
- Ruby on Rails application that allows for development of plugins
- Enterprise edition provides managed service for private code repositories and CI workflows
- Provides features such as activity streams, file browser wiki.
- Prides itself on code review features.

HEROKU

- Supports Ruby, Java, Node.js, Python, PHP, Clojure, and Scala
- For each programming language, project containers contain set of useful tools like Setuptools, Pip, and Virtualenv for Python
- Utilizes Git repository as mechanism for updating containers to the latest developed software – rebuilding and deploying for immediate use automatically
- Ability to execute commands, even start shells, on containers.

OPENSHIFT

- RedHat PaaS offering for cloud, enterprise, or community development
- Uses Git repository for code management and automatic build and deployment
- Support for projects using Java, PHP, Python, Node.js. MongoDB, MySQL, etc.
- Database, logging, and other service hosting at pricing from free up to a cost based on usage.
- Ability to log onto container running project.

GOOGLE APPENGINE

- Google PaaS offering with lots of community support
- Support for Python, Java, PHP, and Go
- Pricing structured by number of instances, network communications, data storage, and SSL
- Tools for running locally and then pushing versions to AppEngine
- Ability to run different versions of same product and roll-back to prior versions
- Easy access to a variety of application logs via web console

AWS ELASTIC BEANSTALK

- Support for Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker
- Ability to user servers like Apache, Nginx, Passenger, and IIS
- Multiple concurrent application versions and rollback features
- Access data storage services from Amazon, Microsoft SQLServer, Oracle, etc.
- CloudWatch tool for monitoring application
- Amazon SNS support for sending emails when application status changes

UNITY3D

- Ecosystem for building and deploying multiplatform games
- Develop games completely within environment with tools
- Incorporates features for maintaining and expanding user base via ads, social sharing, etc.
- Marketplace for buying/selling of game assets
- Incorporation of ads within games
- Monitor user and usage statistics

Module 3: Testing

Web Service Testing

- SoapUI
- SoapSonar
- Applause

Unit Testing

- SauceLabs
- BlazeMeter
- JUnit
- Nose
- xUnit

Load/Performance Testing

- Neotys
- Apache Jmeter
- Loadtesting tool.com
- TestStudio
- LoadComplete

Code Coverage (per Wikipedia)

- In computer science, code coverage is a measure used to describe the degree to which the source code of a program is tested by a particular test suite.

- A program with high code coverage has been more thoroughly tested and has a lower chance of containing software bugs than a program with low code coverage. Many different metrics can be used to calculate code coverage; some of the most basic are the percent of program subroutines and the percent of program statements called during execution of the test suite.

- Code coverage was among the first methods invented for systematic software testing. The first published reference was by Miller and Maloney in *Communications of the ACM* in 1963.

Module 4: Configuration Management

- Operating System CM
  - Ansible
  - Puppet
  - Chef
  - Vagrant
  - Salt
- Programming language CM
  - virtualenv
  - gvm
  - rvm

Module 5: Deployment Tools

WPKG
RPM
Leroy
PDQ
Serena
AWS OpsWorks

Module 6: Management and Alerting

CopperEgg

NewRelic

ManageEngine Applications Manager

Zenoss Cloud Monitoring

AppDynamics

FrameFlow

AmazonCloudWatch

Module 7: Caches and Queues

Caches

SuperCache Express
Memcached
Redis
SQLite
GridGain
HazelCast

Queues

RabbitMQ
ActiveMQ
ZeroMQ
Apache Kafka
HornetQ
IronMQ
Celery Task Queue
AWS SQS
Microsoft Message Queuing

Module 8: Application and Web Servers

- Java
  - GlassFish
  - Jetty
  - Tomcat
  - WebLogic
  - WildFly
- .Net
  - Windows Server
  - IIS
- Python
  - Flask
  - Django
- Miscellaneous
  - Apache HTTP Server
  - NGINX
  - Node.js
  - Ruby on Rails

Module 9: Security

Qualys
WhiteHat Security
Okta
Zscaler
CipherCloud
DocTrackr

Centrify
Vaultive
Tinfoil Security
Fortify on Demand
AWS IAM
Microsoft Azure Trust Center

Security

- Qualys
- WhiteHat Security
- Okta
- Zscaler
- CipherCloud
- DocTrackr
- Centrify
- Vaultive
- Tinfoil Security
- Fortify on Demand
- AWS IAM
- Microsoft Azure Trust Center

Qualys

- Performs system-level vulnerability analysis/management
- Ability to scan for available web applications, check for vulnerabilities, and tries to block attacks via this vector.
- New features aim to detect and protect against advanced persistent threats via recognizing incoming and existing malware
- Monitors assets as they come online for standards compliance

WhiteHat Security

- Software-as-a-Service that provides a vulnerability management for web applications
- Applies continuous rigorous assessment of applications providing 100% vulnerability identification
- Provides three stages of security testing, pre-production, production, and mobile application.

Okta

- Has three main product lines: mobility management, identity management, and identity platform
- Mobility Management – Integrated identity and mobility management solution delivered entirely in the cloud
- Identity Management – Provide ability to create and manage rich user profiles.
- Identity Platform – Authenticate and manage users in the cloud, based on Active Directory. Allows applications to delegate authentication and synchronize profiles in real-time across multiple platforms

ZScaler

- Provides SSL/HTTPS capability to web applications
- Provides access to security appliances that aim to recognize/remediate zero-day attacks and advanced persistent threats

- Risk assessment of current security and compliance infrastructure with recommendations for closing any gaps.

- Zscaler’s global Security as a Service platform acts like a series of global check posts in the cloud between employees and the Internet.

CloudCipher

- CipherCloud provides comprehensive visibility and control over your data as it goes from your enterprise to any location in the cloud.
- By providing a control point for data going to and from the cloud, CipherCloud makes it easy to ensure data privacy,data residency and regulatory compliance, prevent data leaks, encrypt or tokenize sensitive data and get unrivaled visibility into cloud activity.
- Monitor data coming into and out of enterprise, focusing on data protection and reducing malware transmission

DocTrackr

- Information Rights Management (IRM) approach attempts to provide seamless protection to data at rest, in transit, or in-use
- Rights and permissions may be updated and revoked even after a file has been downloaded
- Each document is assigned a unique encryption key, and asymmetric cryptography is used in the distribution of the key to each user
- Maintains a complete audit trail of all files and user interactions, aiding organizations in regulatory compliance.

Centrify

- Identity Management and Auditing for Big Data – leverages an organization's existing Active Directory infrastructure to deliver access control, privilege management and user-level auditing across Hadoop clusters, nodes and services.
- Integrates with Thousands of Apps and Systems such as Sharepoint, Office 365, JBoss, and Salesforce
- Centrify Identity Service provides Single Sign-on for Every App
- Multi-Factor Authentication (MFA) and Policy capabilities including context-aware, step-up MFA based on per-app policies

Vaultive

- Configurable encryption engine secures and protects data in almost any cloud-based service
- Operates as a network-level proxy, encrypting data at the edge of your corporate network, before sending it to the cloud service provider
- Server-side functionality, such as search, sort and indexing is maintained without ever decrypting the data.

Tinfoil Security

- Heroku add-on that detects and manages vulnerabilities in Heroku applications
- Provides simple dashboard into application security issues
- Tiered pricing structure allows for configurable scanning parameters such as scan frequency, max number of pages to be scanned, and scan types

Fortify on Demand

- Scans for Open Web Application Security Project (OWASP) top 10 vulnerabilities
- Attempts to identify issues with applications early to reduce the possibility of attacks
- Penetration testing performed by humans

AWS Identity and Access Management (IAM)

- By creating users through IAM, can manage user credentials (passwords, keys, and multi-factor authentication) to AWS resources
- Roles can be assigned to IAM-based users, to allow control of what actions users can perform
- Allows for incorporation of users that already exist in your enterprise, removing the need to create IAM for these users

Microsoft Azure Trust Center

- Integrated deployment systems manage the distribution and installation of security Datacenters are physically constructed, managed, and monitored to shelter data and services from unauthorized access as well as environmental threats.
- Security is monitored with the aid of centralized monitoring, correlation, and analysis systems that manage the large amount of information generated by devices within the environment and providing timely alerts & patches.
- Microsoft Antimalware is built-in to Cloud Services and can be enabled for Virtual Machines to help identify and remove viruses, spyware and other malicious software and provide real time protection.

Module 10: Dashboards and Supervision

iDashboards
ServerDensity
Aternity
Boundary
Compuware APM

How do we learn so many new tools?

- Develop technical modules
  - self-study
  - examples
  - To achieve standards of proficiency
- Adopt one or two into the teaching of each course as a practical component

Success metrics – (Nonlinear Digital)

- **#1. Mean time to recovery/repair (MTTR)** (Incident report to resolution)
  - Measures responsiveness of the team
  - Indicates capability to resolve and deploy solutions.
- **#2. Lead time** - Time from start of development to production. Teams optimizing this metric
  - Attempt smaller chunks of work
  - Optimize integration of the testing process
  - Resulting in shorten overall time to deployment.
- **#3. Percentage of successful deployments** (Customer Accepted)
  - Is not just about avoiding outage
  - Goal is to maintain positive reactions in the customer base.
  - Goal is to optimize smaller, less risky changes more often.
- **#4. Projects completed per quarter**
  - Reports success to management on the team's ability to execute.
  - Provide confidence that the investment in the team has generated results.
  - Promotes teams to try to optimize for smaller batches to get more projects done in a shorter amount of time.
  - Leads to delivery processes that support multiple projects at once to avoid the classic "project queue" in testing environments.

DevOps lifecycle – (New Relic)

- Create new code changes [Business Requirement]
- Check in code (GIT)
- Pull code changes, check latest patches, and build (Tools)
- Test (continuous integration server to generate builds & arrange releases):
  - individual models
  - integration tests
  - user acceptance tests.
- Store artifacts and update the build repository (repository for storing artifacts, results & releases) (Staging/Configuration Management)
- Deploy and release (release automation product to deploy apps) (Release Management
  - Configure environment
  - Update databases
  - Update apps
  - Push to users
- Application & Network Performance Monitoring (preventive safeguard)
- Repeat

Recent Studies Find

- Everyone wants DevOps – yet definitions, standards, and guidelines are lacking
- Ops needs stability first while Developers want innovation
- Cultural barriers between developers and operations are the greatest obstacles