



Essential insight for leaders
of the digital economy.

Docker Orchestration

@cote – 13th Nov 2014

451 Research is an information technology research & advisory company

Research & Data

Advisory Services

Events

Founded in 2000

210+ employees, including over 100 analysts

1,000+ clients: Technology & Service providers, corporate advisory, finance, professional services, and IT decision makers

10,000+ senior IT professionals in our research community

Over 52 million data points each quarter

4,500+ reports published each year covering 2,000+ innovative technology & service providers

Headquartered in New York City with offices in London, Boston, San Francisco, and Washington D.C.

451 Research and its sister company Uptime Institute comprise the two divisions of The 451 Group



Michael Coté

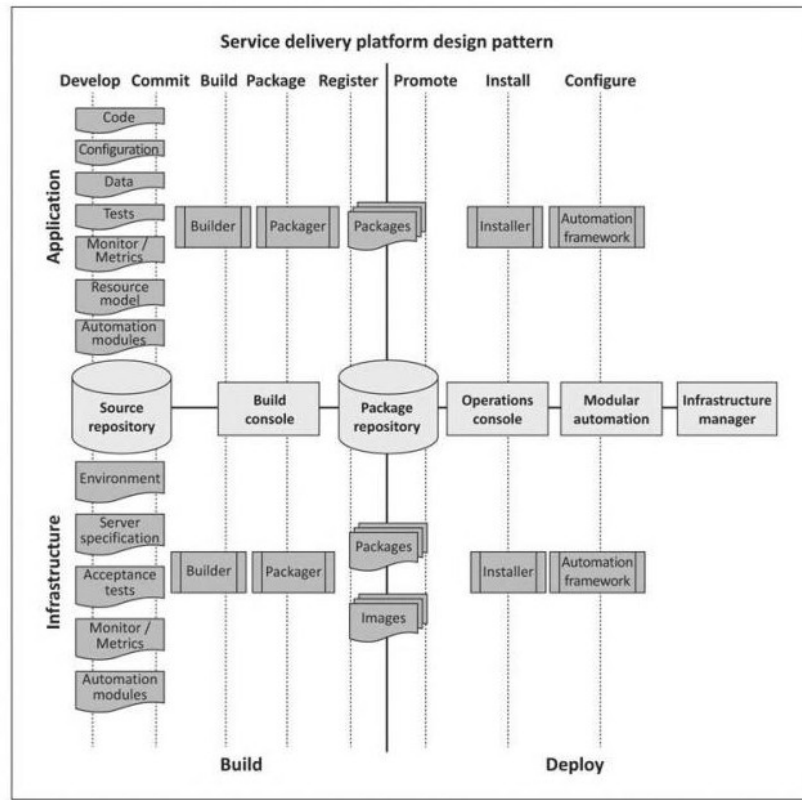
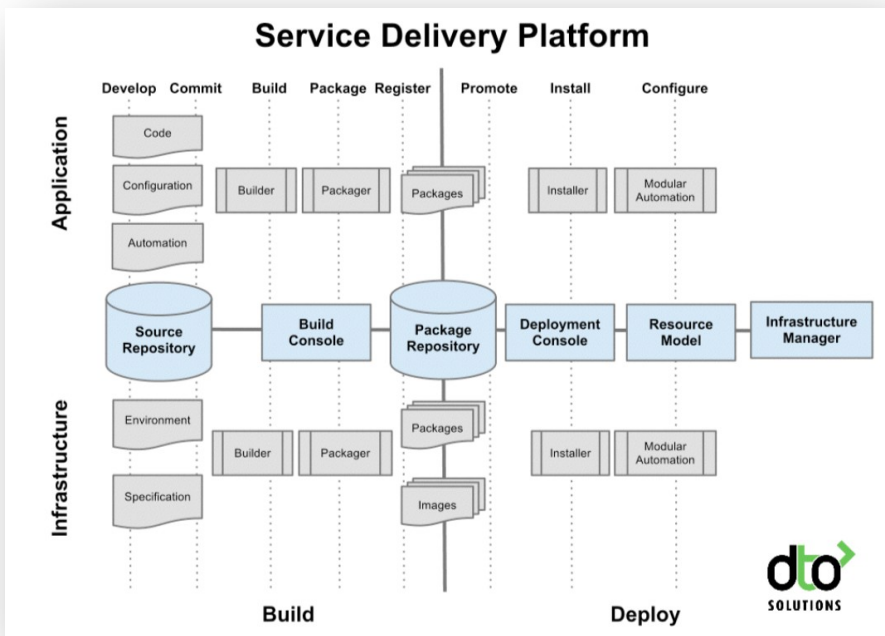
- Research Director of Infrastructure Software, 451 Research
- Focusing on developers and cloud
- I've worked in development at startups and large ISVs, as an analyst, and in corporate strategy & M&A.

- cote@451research.com
- [@cote](#)
- <http://cote.io>



What's
driving
demand?


Custom written software is a business' most valuable IT asset, cloud and DevOps are helping open up that value to a wider audience



PaaS has so far been anemic from a charts perspective

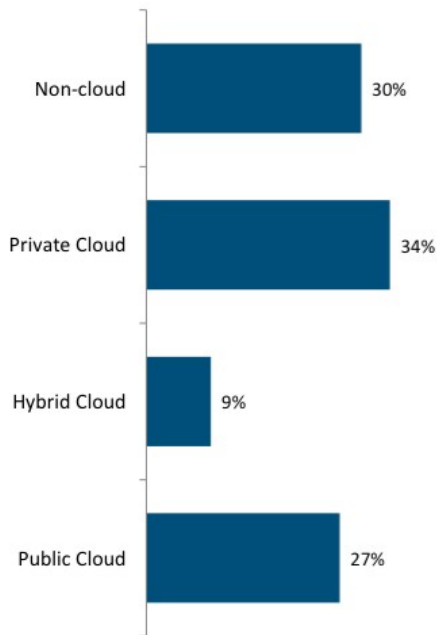
CaaS - comparing cloud markets (\$bn)

(Note: OpenStack is sub-set.)



Over the next two years, enterprises are looking more to private than public cloud

Primary Deployment Method



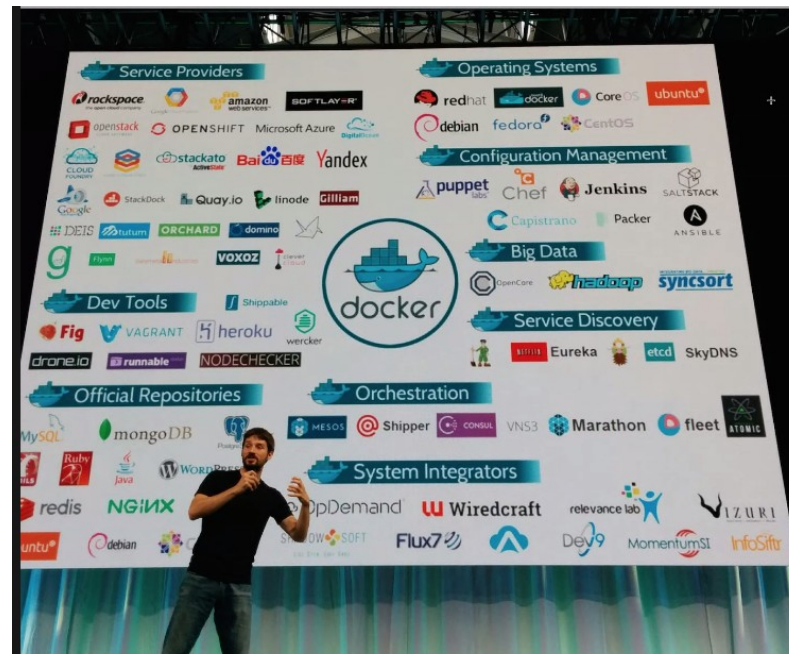

Leading Reasons for Deployment Choice



Notes: Left Chart: Q. For each of the following categories of workload/business functions, what is your primary deployment method likely to be in the next two years (internal non-cloud; internal, on-premises private cloud; on-premises, hosted private cloud; off-premises, non-cloud hosting; off-premises, hosted private cloud; hybrid cloud; external public cloud; or SaaS)? n=863. Right Chart: Q. What are the reasons for that selection? n=645.

Source: [“TheInfoPro Cloud Computing Wave 7: Cloud Computing Metrics,”](#) 451 Research, 10 Sep 2014.

Teams are eagerly looking to new technologies to build their clouds



Notes: Automated Server Provisioning usage and planned time series, 2H '13, n=118; 1H '14, n=162.

Source: "[TheInfoPro In-depth Technology Roadmaps - Cloud Computing Wave 7](#)," 451 Research, 22 Aug 2014; Docker logo slide from Alex Williams.



Docker Orchestration

Emerging requirements for Docker orchestration

- 1) Cluster/fleet management
- 2) Configuration management & Automation
- 3) Heterogeneous platform support
- 4) Baby and bathwater-ing ITSM



Cluster/fleet management

Operate in terms of multiple nodes, not single nodes

Configuring a single Docker node is mind-blowingly easy, doing it over 50 or 100 nodes gets to be tedious, esp. if you want to continually be turning over builds. Pets vs. cattle and all that.

Configuration management & automation

Application modeling that describes the layout and configuration of various components

This is an old ITSM notion, "service modeling," which got bogged down in drag and drop fantasy (just like UML). You need to model what all the different components are and how they fit together.

Basic CRUD

Creating nodes, updating nodes, restarting them as needed. You want more than just modeling what a node looks like, you want you orchestrator to actually do something.

Configuration management & automation, cont.

Separating configuration from basic state

Easily modify configuration without having to change too much about each node/image, like changing port numbers easily without rebuilding the entire node.

Ensure proper configuration passing across nodes

Passing server names and ports to servers, handing out credentials, wiring in service directories, etc.

Heterogeneous platform support

"Hybrid cloud" or "multi-cloud" support

Support for different infrastructure, bare-metal, to plain old virtualization, to multiple clouds. Also useful for different stages in the delivery pipeline (dev, QA, staging, prod).

Baby and bath-watering ITSM

Asset database to track all your cows, aka, “source of truth”

This starts getting into “enterprise” needs, but it’s handy even if you aren’t tweedy. You need to know what you have out in the wild and quickly locate it when things go wrong.

Capacity management and adjustment of resources


Not only monitoring if you’re over (or under!) capacity, but actually going back to your CRUD operations to make adjustments on your nodes. This is also where keeping configuration separate from node state is handy: you could increase memory, keeping the same configuration, for example, without having to rebuild or swap out nodes.

Ease of use, low cost

Ease of use, and esp. low cost

Otherwise, why not just use a full on PaaS? This is always easy to forget, but it's sort of the point of all of this. Ask yourself, is this easy and quick to use? If it's not, something is wrong.

Good job, you've gotten a blinking cursor up! Now, the hard work can begin.... Coffee is for coders!



Thank You!

@cote – <http://cote.io>



NEW YORK
LONDON
BOSTON
WASHINGTON, D.C.
SAN FRANCISCO