



# Moving Apache Ignite into Production: Best Practices for Deploying Apache Ignite in the Cloud

Greg Stachnick

[greg.stachnick@gridgain.com](mailto:greg.stachnick@gridgain.com)

@gstachni



# Agenda



- Define Terms
- Deployment Options
- Best Practices



Photo by [Jon Tyson](#) on [Unsplash](#)



# Ignite & GridGain In-Memory Computing Platform



Monitoring & Management

Security & Auditing

Segmentation Protection

Data Center Replication

Rolling Upgrades

Connectors (Kafka, GoldenGate, Hadoop)

Application Layer



Web-Scale Apps



IoT



Mobile Apps



Social Media

Machine and Deep Learning

Key-Value

SQL

Transactions

Messaging

Streaming

Events

Compute Grid

Service Grid

In-Memory Data Store

Data Layer



RDBMS



GridGain



NoSQL



Hadoop

Full, Incremental, Continuous Backups

Network Backups

Point-in-Time Recovery

Heterogeneous Recovery



**GridGain** 

+

 **CLOUD NATIVE**  
COMPUTING FOUNDATION

 **LINUX**  
FOUNDATION

# What is “The Cloud”?





# According to Salesforce



1. Flexibility
2. Disaster Recovery
3. Automated Software Updates
4. Capital-expenditure Free
5. Increase Collaboration
6. Work From Anywhere
7. Document Control
8. Security
9. Competitiveness
10. Environmentally Friendly

<https://www.salesforce.com/uk/blog/2015/11/why-move-to-the-cloud-10-benefits-of-cloud-computing.html>

# According to IBM



- Scalability
- Storage Options
- Control Choices
- Tool Selection
- Security Features
- Accessibility
- Speed to Market
- Data Security
- Savings on Equipment
- Pay Structure
- Streamlined Work
- Regular Updates
- Collaboration
- Competitive Edge

<https://www.ibm.com/cloud/learn/benefits-of-cloud-computing>

# Elastic Scalability





# Flexibility



Photo by [Yannes Kiefer](#) on [Unsplash](#)

# It's not in your data center



Photo by [Taylor Vick](#) on [Unsplash](#)



# Best Practice 1: Tooling



Photo by [Lachlan Donald](#) on [Unsplash](#)



# Docker is...



“A container is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another”



By dotCloud, Inc. - File:Docker (container engine) logo.png,  
Apache License 2.0, <https://commons.wikimedia.org/w/index.php?curid=52332268>

# Kubernetes is...



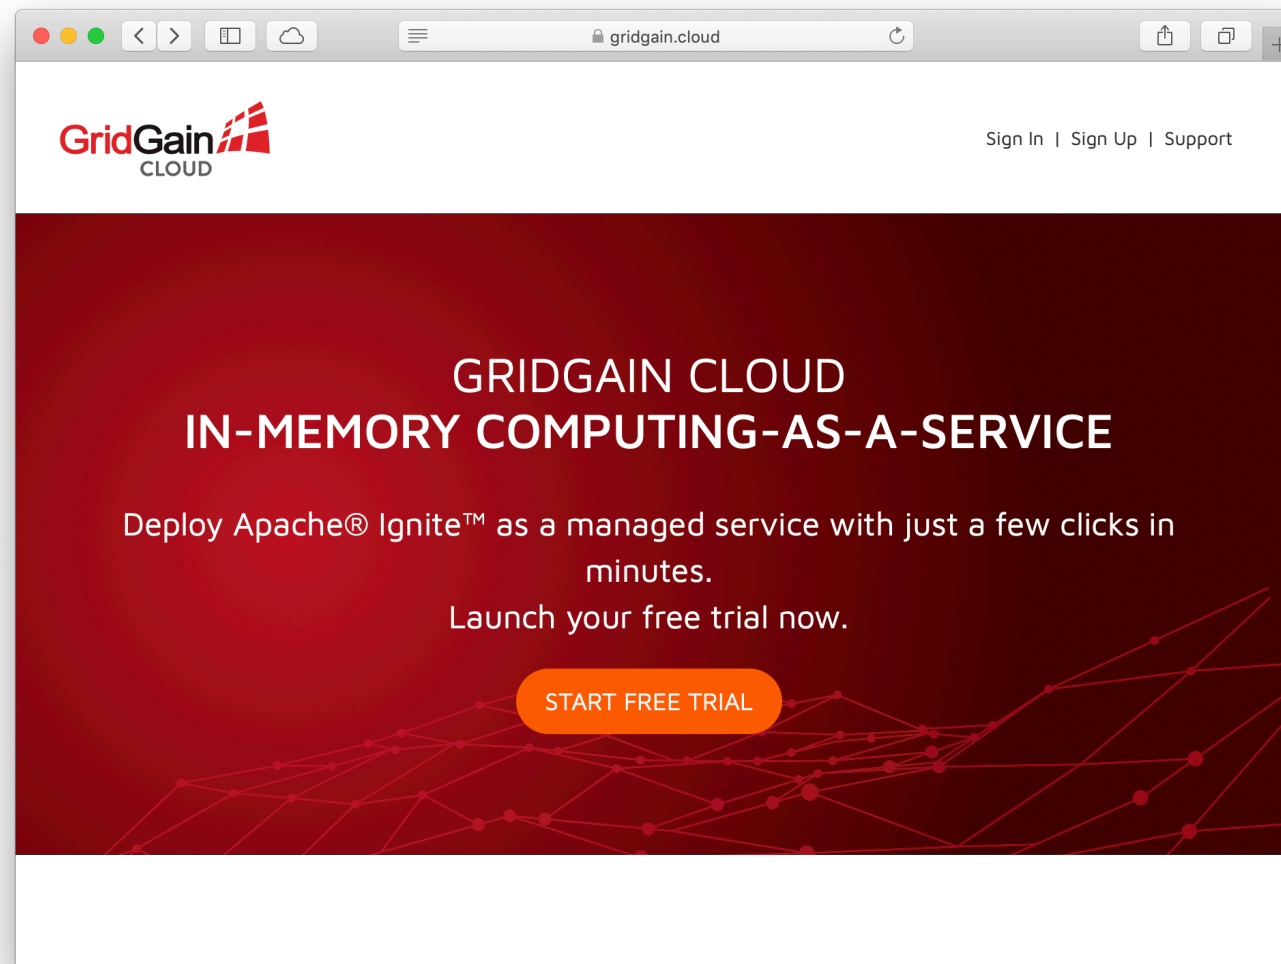
“...an open-source system for automating deployment, scaling, and management of containerized applications.”



# SaaS: GridGain Cloud



- Fully Managed In-Memory Computing Built on Apache Ignite
- Web Console
- Data access
  - REST
  - JDBC/ODBC
  - Ignite thin-clients
- “Up and running in minutes”







- EC2
- ECS
- EKS
- On-demand
- Spot instances
- Reserved instances
- Dedicated hosts
- Fixed performance
- Burstable
- Cluster Networking
- Intel
- ARM
- General purpose (7 options)
- Compute (3 options)
- Memory optimized (7 options)
- Accelerated (4 options)
- Storage optimized (4 options)



- Virtual Machines
- Virtual Machine Scale Sets
- Azure Kubernetes Service
- Container Instances

# Best Practice 2: Memory Sizing



- For data-grid use cases, aim for the “big memory” options
- Consider Replication
  - Better to have four 256Gb nodes than one 1Tb node
- Use the Sizing Calculator (<https://apacheignite.readme.io/docs/capacity-planning>)




Photo by [Franck V.](#) on [Unsplash](#)

# Best Practice 3: Disk Storage



- Trade offs
  - EBS vs Instance Storage Volumes
- AWS EBS – IOPS SSD (io1) volumes
- Azure – Standard or Premium SSD
- Use StatefulSets in Kubernetes



Photo by [Vincent Botta](#) on [Unsplash](#)

# Best Practice 4: Images



- Use custom images or containerization!
- Make it easy to create / kill new GridGain nodes
- Automate as much as possible



[This Photo](#) by Unknown Author is licensed under [CC BY-SA](#)



# Best Practice 5: Performance



- Predictable versus maximum
  - Reserved
  - Spot
- Scale automatically
  - “Horizontal Pod Autoscaler” with Kubernetes
  - EC2 Auto Scaling in AWS
  - Scale Sets in Azure
  - But remember licensing



Photo by [Alessio Lin](#) on [Unsplash](#)

# Best Practice 6: Security



- TLS/SSL between nodes
- Disk encryption
- Firewall ports
- Use GridGain security options for authentication, authorization and auditing




Photo by [Victor Garcia](#) on [Unsplash](#)

# Best Practice 7: Data Location



- Where is your data?
- Who has access to data?
  - Synchronization Layer
  - Change Data Capture
- How do you get it to “the cloud”?
  - Deltas versus full extracts
  - Migrate everything?



Checkout [\*In-Memory Computing Best Practices: Adding Speed and Scale to Existing Applications\*](#)

Photo by [Mika Baumeister](#) on [Unsplash](#)




# Best Practice 8: Migration

- Use Data Center Replication



Photo by [Jason Hafso](#) on [Unsplash](#)

# Best Practice 8: Migration





# What have we learned?



- The cloud *is* different from on-prem
- Best practices
  - Tooling
  - Memory sizing
  - Disk storage
  - Use custom images
  - Performance
  - Security
  - Data location
  - Migration




Photo by [Elijah Hail](#) on [Unsplash](#)

# GridGain Resources



- Webinars
  - Visit <https://www.gridgain.com/resources/webinars>
  - Visit <https://www.imcsummit.org/>
- White Papers
  - Visit <https://www.gridgain.com/resources/papers>
- Videos
  - Visit <https://www.gridgain.com/resources/videos>
- Sign Up for The In-Memory Computing Summit (Nov 13-14)
  - Visit <https://www.imcsummit.org/>
- Free 30-Day Ultimate, Enterprise or Professional Edition Trial
  - Visit <https://www.gridgain.com/resources/download>

# Thank you

Greg Stachnick, Director of Product Management

[greg.stachnick@gridgain.com](mailto:greg.stachnick@gridgain.com)

@gstachni

