



GridGain Nebula Managed Service: How We Deploy Gridgain and Apache Ignite in Clouds

Andrei Alexandrov

About Me



Name: Andrei Alexandrov

Company: GridGain since 2017

Team: Customer Solutions

Email: aealexandrov@gmail.com

Agenda



- What are the benefits of public cloud deployment?
- What issues can arise during public-cloud deployments?
- How does GridGain Nebula managed services resolve these issues?

Three Reasons Why People Prefer Public Clouds



- Simple host management
- Possibilities of flexible host customization
- Support and deployment on remote hardware

aws



Azure



Simple Host Management



How to manage a server:

- Create an account.
- Configure your host (web interface, CLI, programmatic API).
- Start your host.



Possibilities of Flexible Host Customization



To get a solution that meets your goals, you can do any or all of the following:

- Use the host configurations that you require (memory, storage, CPU)
- Use the features that you require (public IP, DNS, Load Balancer, and so on)
- Use managed services for your own solutions (Kafka, Spark, Hadoop, and so on)



What comes with HDInsight?



Apache Hadoop



Apache Spark



Apache Kafka



Apache HBase



Apache Hive LLAP



Apache Storm



Machine Learning

Support and Deployment: Bare Metal Deployment vs Public Cloud Deployment



Bare metal deployment

- Buy your own server.
- Hire server admins for hardware support.
- Hire a DevOps team for deploying and supporting the required software.
- Hire a development team for implementing business logic.



Public cloud deployment

- Buy a subscription for the cloud hosts and get the support.
- Buy managed services for the required software.
- Hire a development team for implementing business logic.
- Make tier hosts available for development and testing.

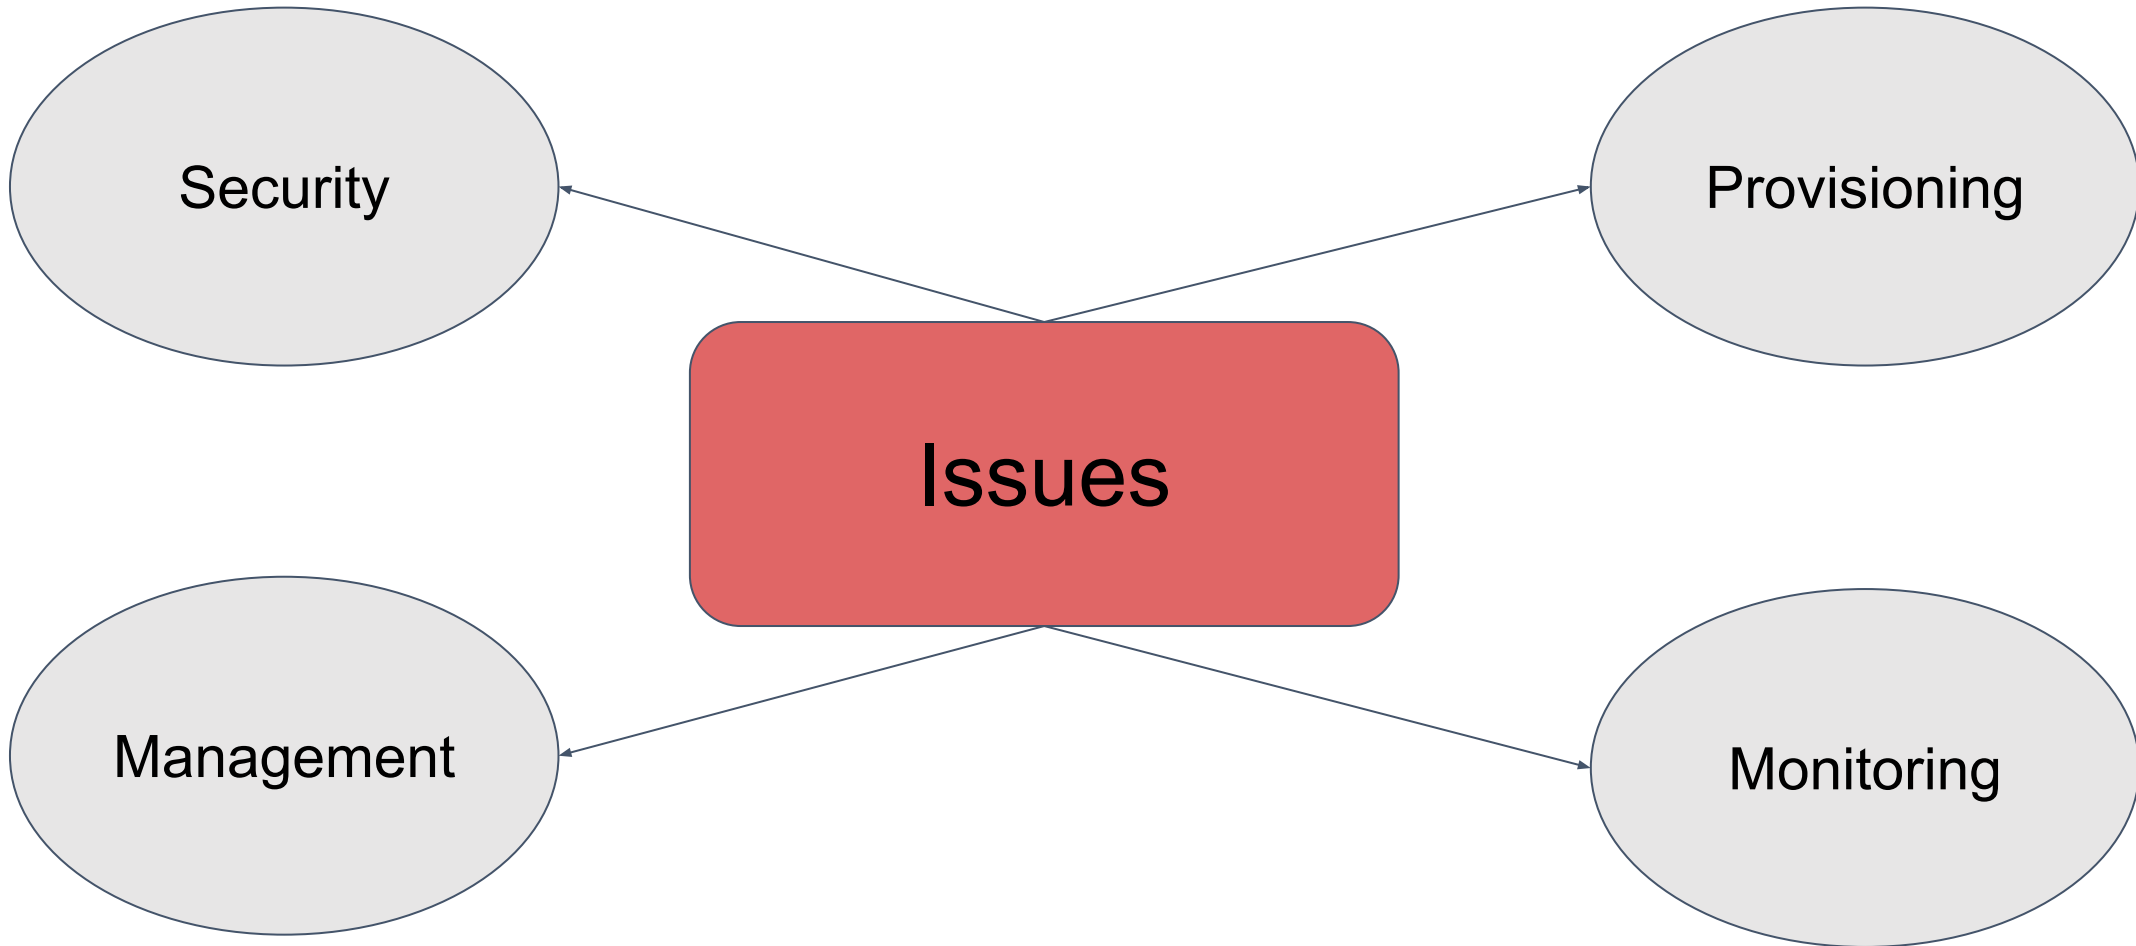
Support and Deployment: Benefits of Using Public Cloud



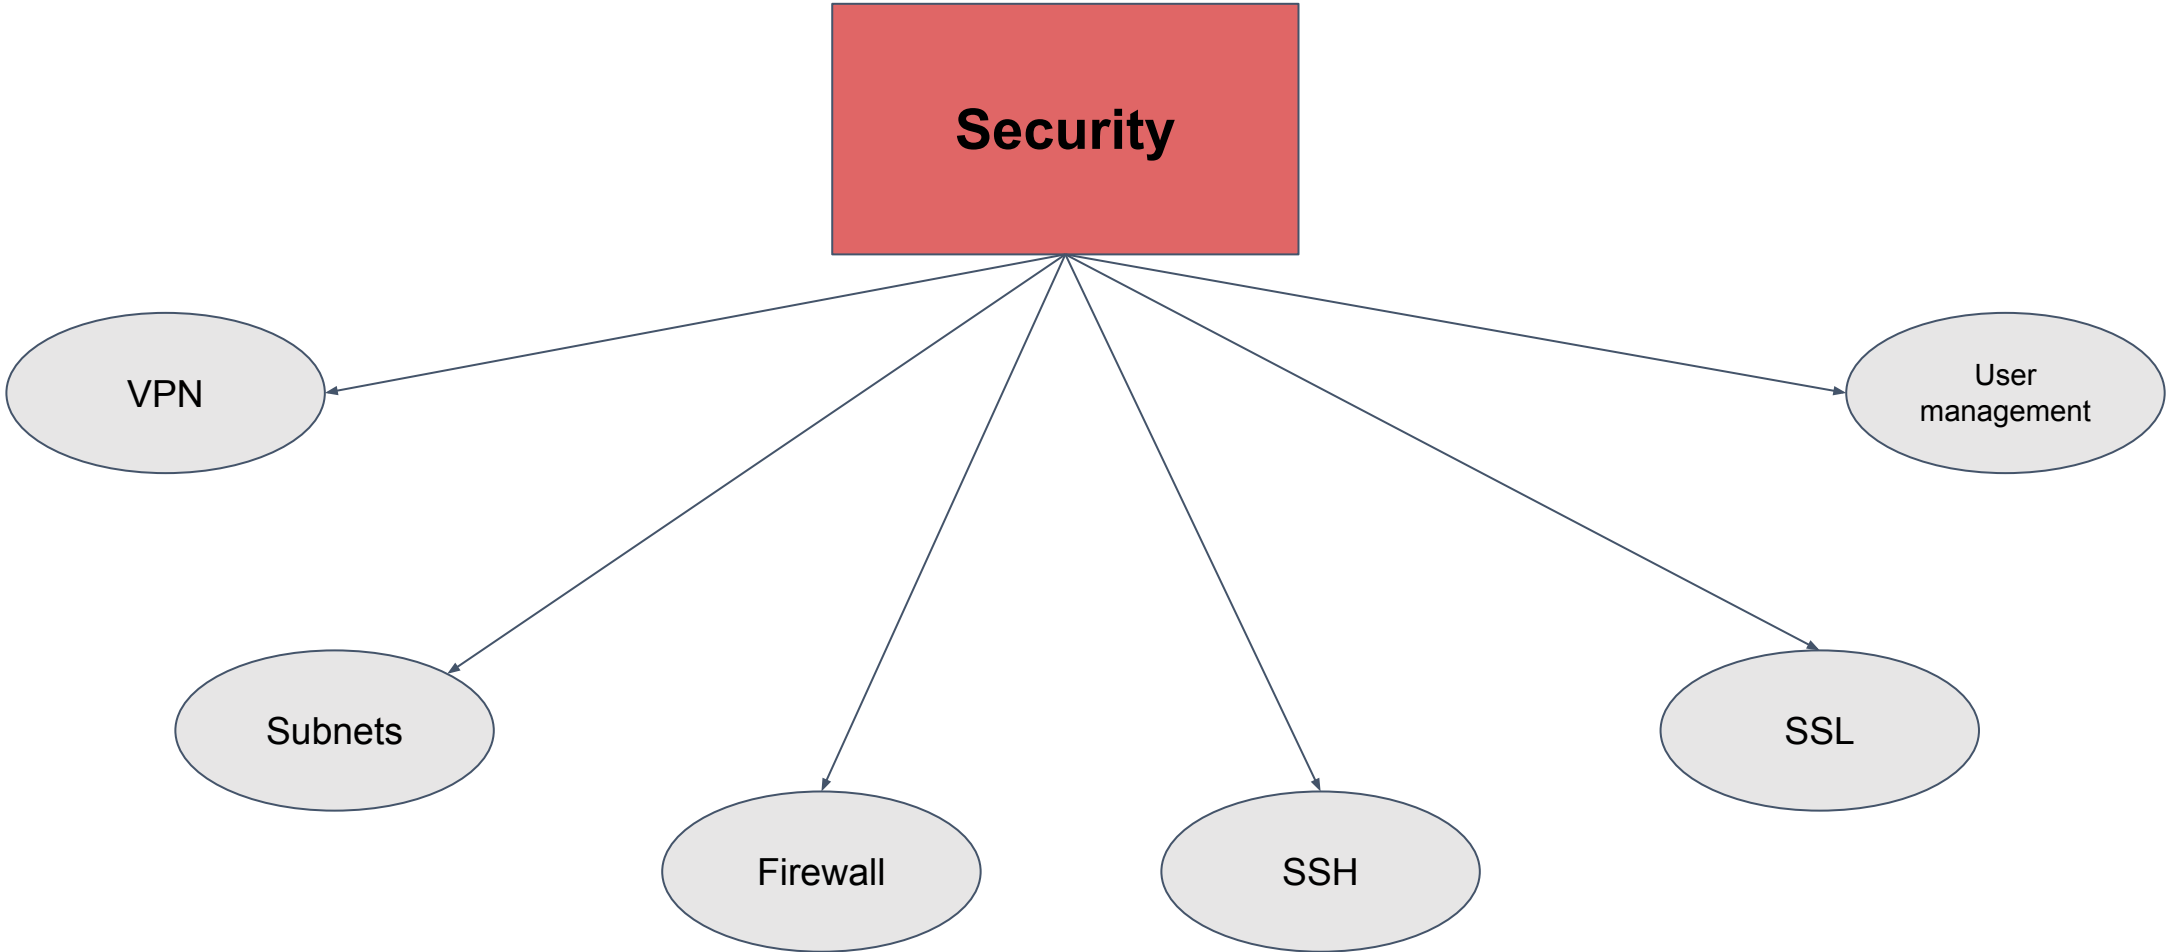
But, how can we save money and time with a public cloud?

- Cloud hardware issues will be resolved faster with cloud support.
- A managed services team can set up software faster than a DevOps team and can provide support.
- Host upgrades do not require modification or replacement of old servers (for example, after five years when new models appear).

Four Main Issues of Public Cloud Deployment



Security

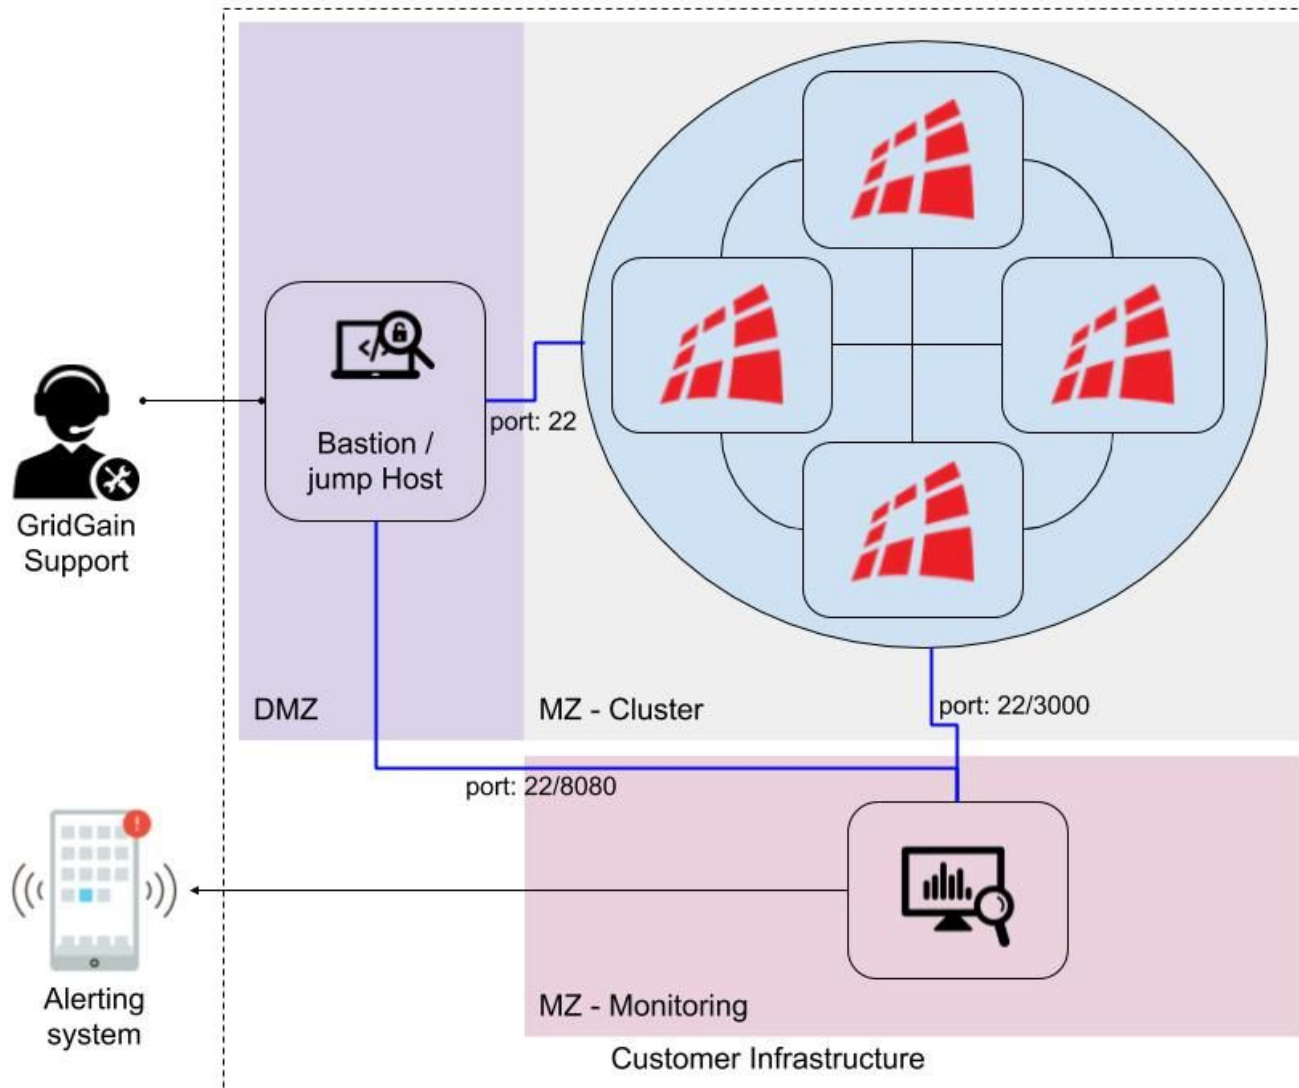


Bastion Host Configuration

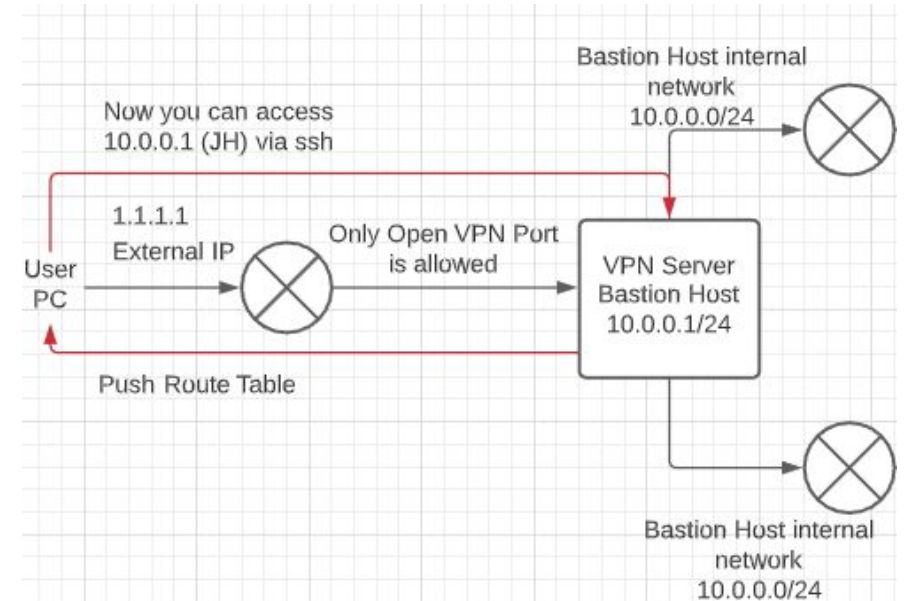


- SSH and VPN configuration
- Firewall configuration between Bastion host and other networks
- Compliance with security standards used in your company

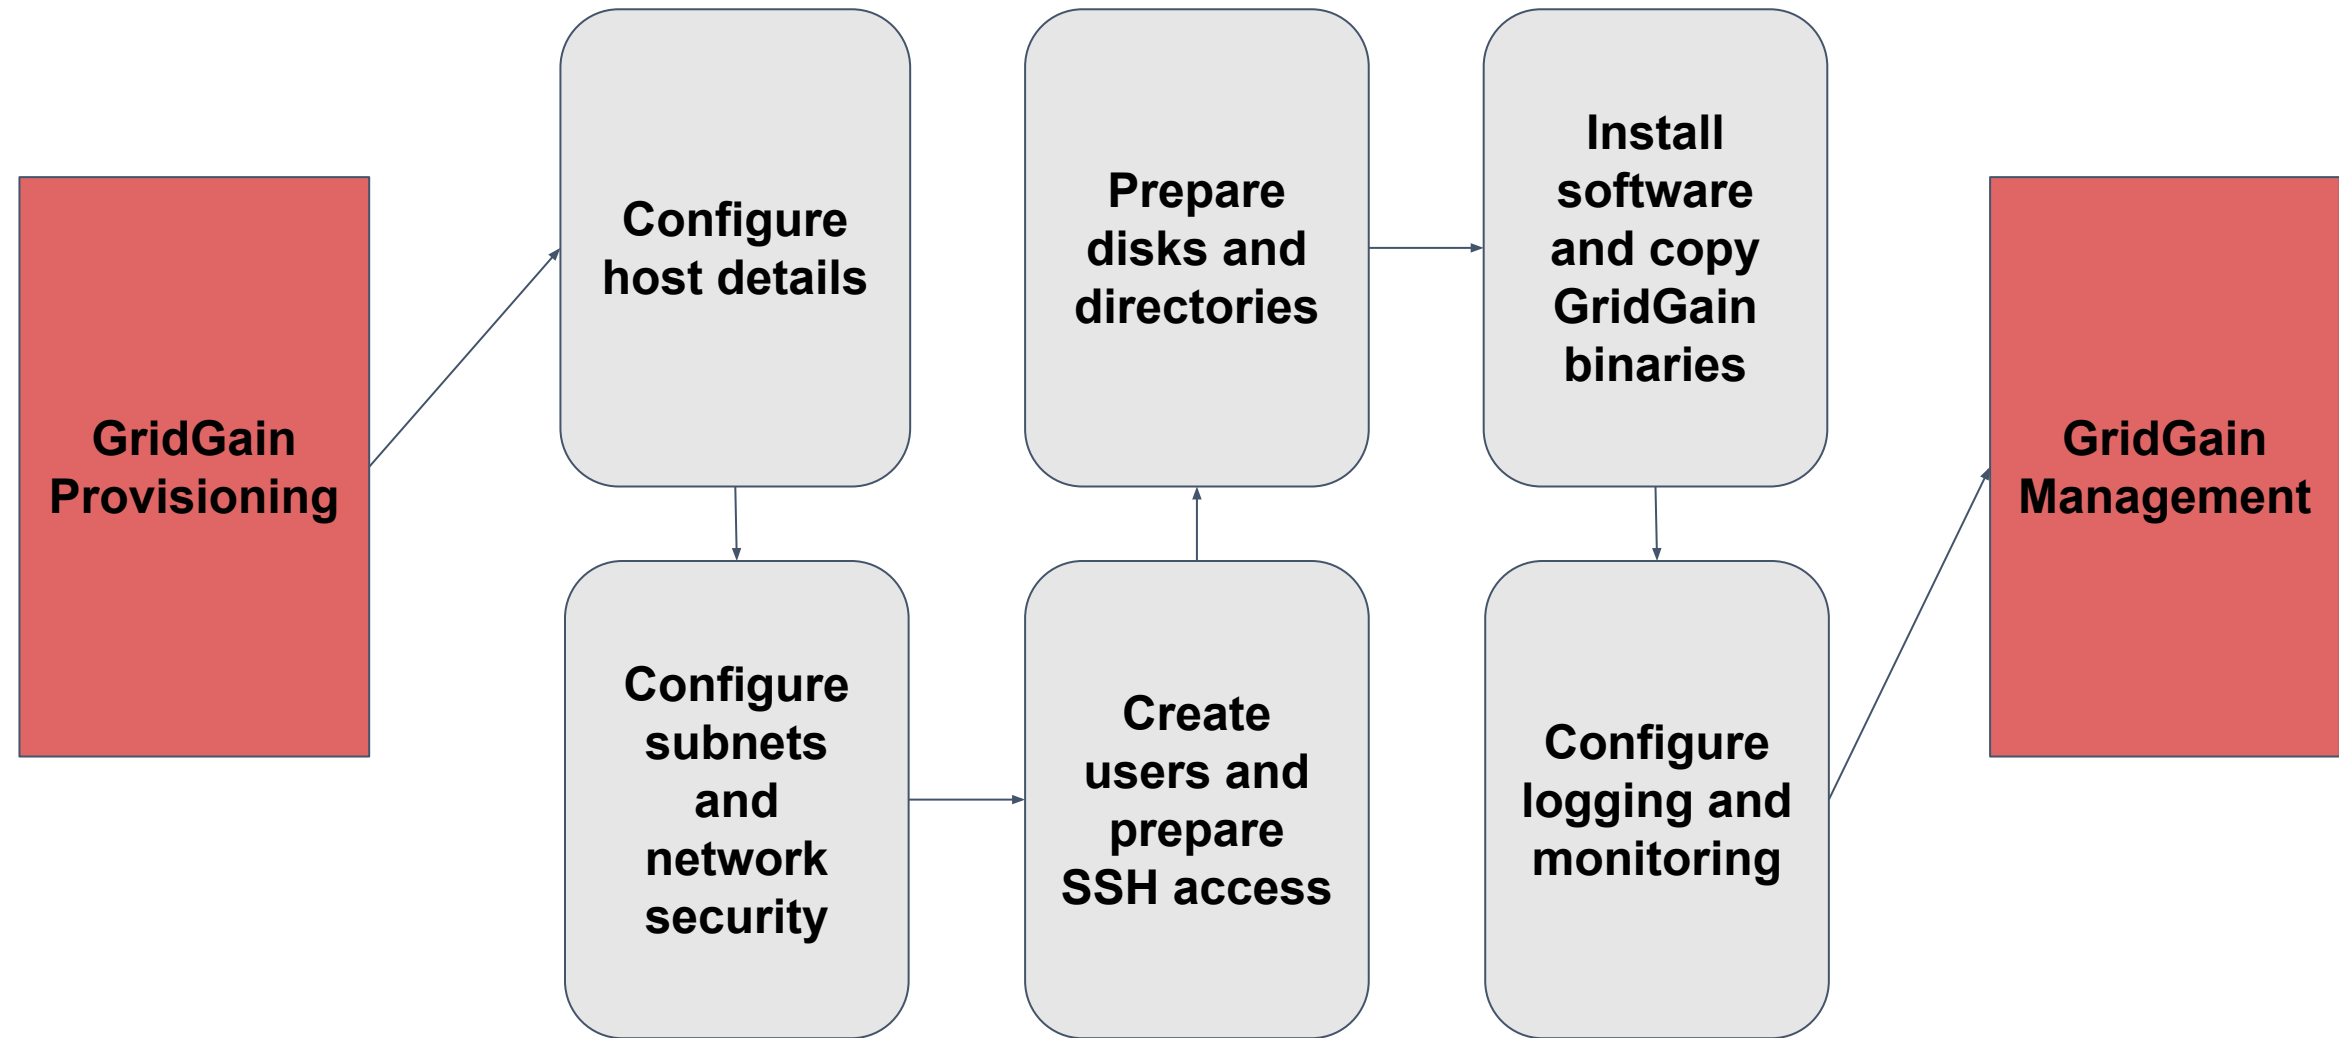
GridGain Nebula Bastion Host Solution



- Three networks (Bastion host, cluster hosts, monitoring hosts)
- SSH access to Bastion host under VPN (see the second diagram)
- Automatization of provisioning (creation and initial configuration of the instances)
- Automatization of the management processes (starting and stopping the GridGain nodes, updating the certificates, getting the logs, and so on)



Provisioning GridGain on a Public Cloud Host



GridGain Management on Public Cloud Host



**GridGain
Management**

**GridGain
stop and
start**

**Rolling
upgrade**

**Updates of
the SSL
certificates**

**And a lot of
other
tasks...**

**Log
retrieval**

**Change
configuring**

**Data
snapshots**

How Can You Implement Provisioning and Management?



With **five** nodes, you can do something like the following:

1. Using a Cloud web console or CLI, configure the hosts and network security.
2. Using SSH, install the packages, create users and directories, and so on.
3. Using SCP, copy the GridGain binaries.
4. Using SSH, start GridGain.
- ...
- N. Repeat steps from 2 to (N-1) **five** times.

But, what if you have **100** steps and **100** hosts?

How Did We Implement It in GridGain Nebula?



ANSIBLE

We automate the provisioning and management processes:

- Each action can be implemented as a separate task (ansible roles).
- Tasks can be grouped into a chain (ansible playbooks).
- The chain of tasks can be executed on a subset of hosts in parallel or sequentially.

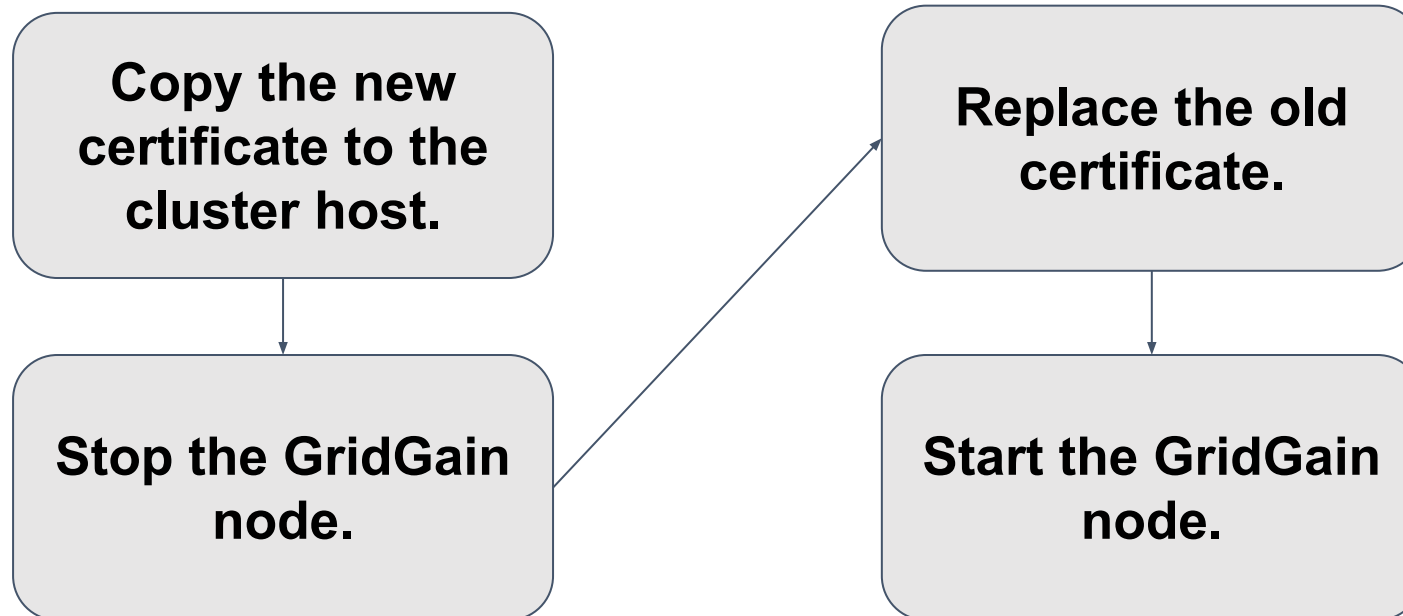
You can read more about Ansible at:

<https://docs.ansible.com/ansible/latest/index.html>

SSL Certificate Renewal Automation



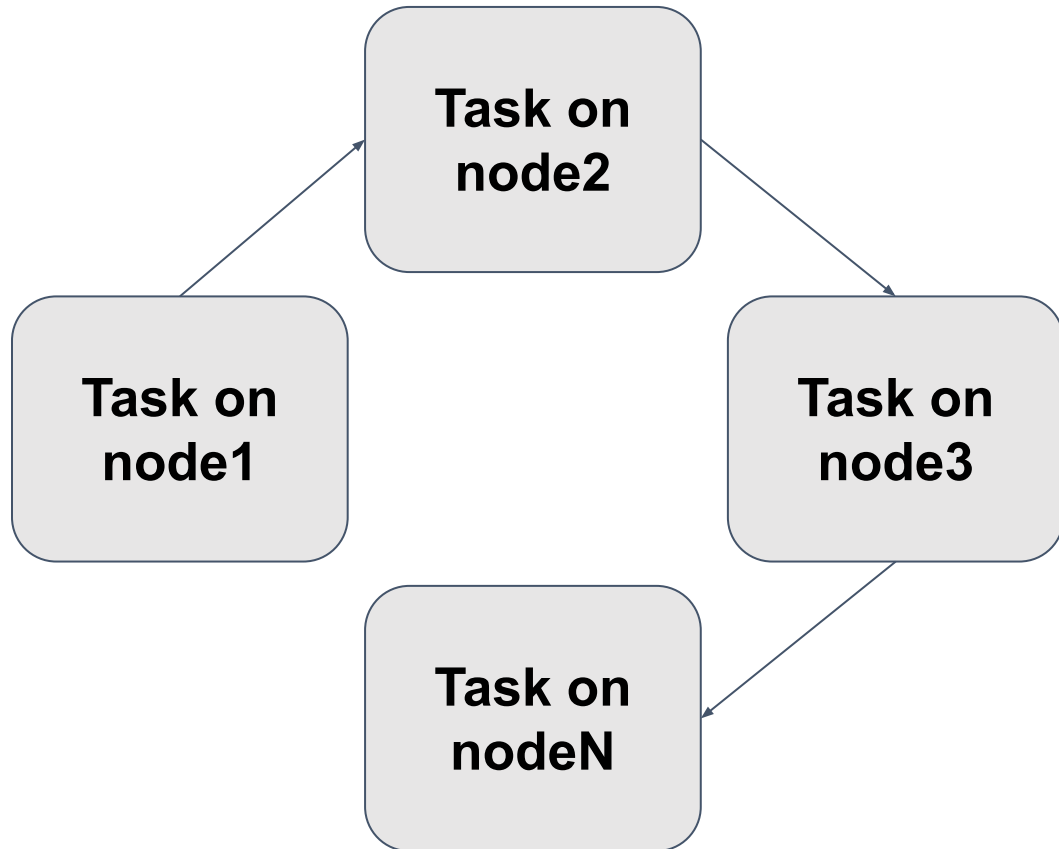
The SSL renewal task for a GridGain node includes the following steps:



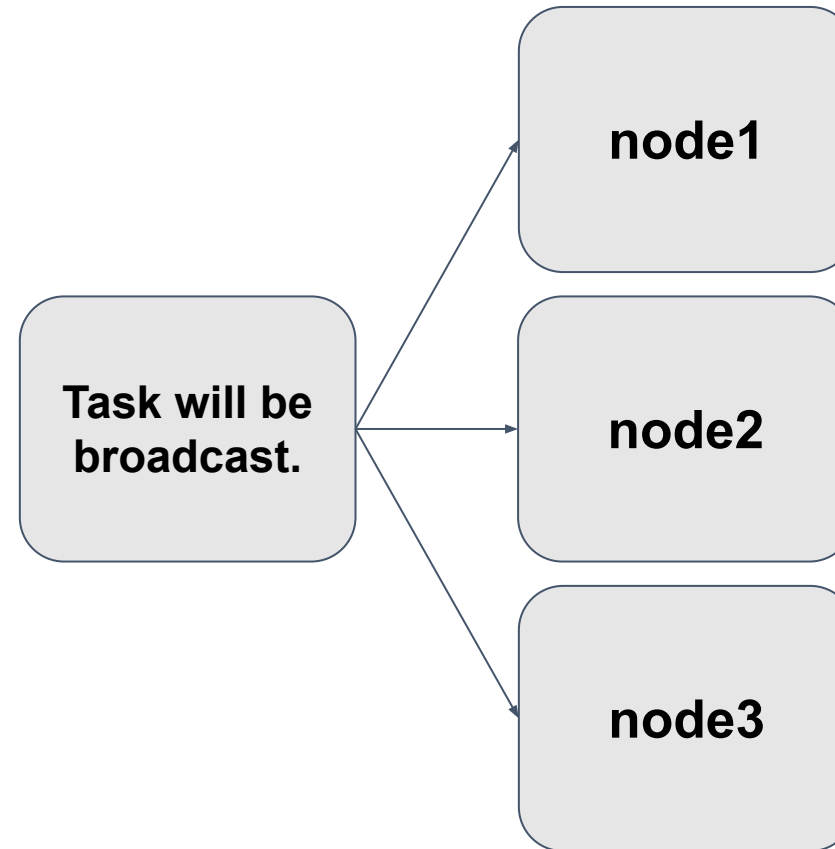
SSL Certificate Renewal Automation



Case 1



Case 2



More Examples



To learn more about how we manage GridGain in the cloud with a live demo, please visit the following webinar in January:

<https://www.gridgain.com/resources/webinars/learn-how-our-managed-services-offering-deploys-gridgain-and-apache-ignite>

Monitoring



We require cluster and system monitoring.

We use GridGain Control Center to monitor the cluster.

<https://www.gridgain.com/docs/control-center/latest/overview>

To monitor the system, by default, we use a Zabbix server:

<https://www.zabbix.com/>

Examples of Monitored Items



System

- Free memory
- Free disk space
- OS audit
- Expiration of the certificates
- Network issues

Cluster

- Memory and CPU usage
- Long running queries
- Transaction tracing
- Rebalance tracking
- Retrieval of information about the data

Summary



GridGain Nebula Managed Service offers all of the following:

- GridGain and Apache Ignite best practices, collected over the years
- Automation of deployment and management processes
- A dedicated support team that can provide fast resolution of any issue
- Extensive monitoring and alarm notification

All that you need do is to implement the business logic.

<https://www.gridgain.com/products/managed-services>