OROCommerce™

How to Build a Reliable and Secure Infrastructure for OroCommerce

DECEMBER 28, 2017

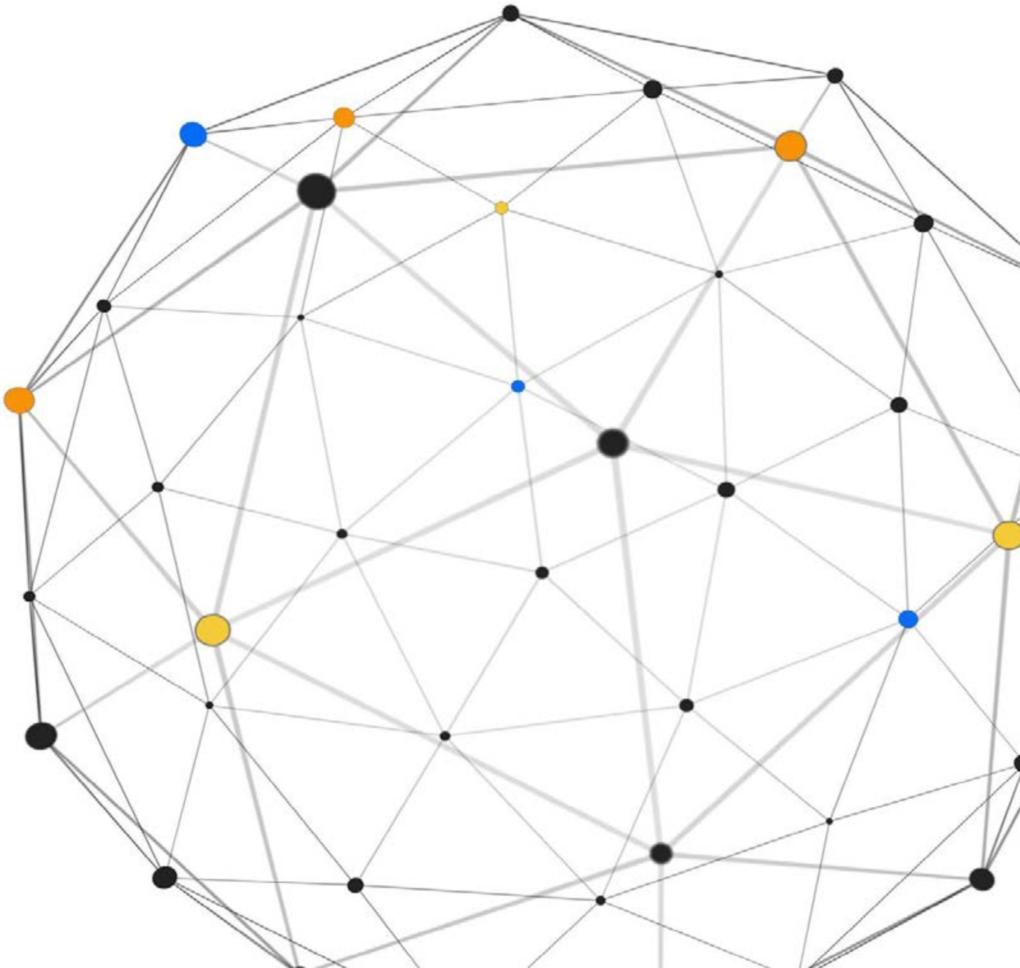

Table of Contents

How to Build a Reliable and Secure Infrastructure for OroCommerce

OroCommerce is the only application built from the ground up for eCommerce targeted towards mid- and large-sized companies. By leveraging an open-source development model and modern architecture, OroCommerce provides the ultimate flexibility to fit any business need and can be integrated to already existing business applications such ERP, PIM, Warehouse Management, etc. In addition, Oro's "deploy anywhere" philosophy allows customers to choose between different deployment models - from an OroCloud SaaS solution to an on-premise installation with prepacked network appliances in popular IaaS platforms, containers and distributives. This means that Oro customers can have a wide variety of platforms and solutions to choose from when setting up an on-premise installation.

While OroCommerce is available both as a SaaS or PaaS solution (which off-loads all support burden for Oro customers), some organizations still prefer to deploy all IT systems in-house or use popular IaaS services like Google Cloud Platform (GCP), Amazon Web Services (AWS) or Microsoft Azure.

This white paper will focus on guidelines and recommendations for building an on-premise infrastructure to host the OroCommerce solution. There are many factors that affect the infrastructure architecture like organizational preferences, budget, and corporate standards just to name a few. Let's review main factors to consider during the infrastructure planning stage.

Things to Consider

Correct planning of the hosting infrastructure is vital to the application's life-cycle and decisions made during this phase will heavily affect the performance of the solution in addition to operational budget and information security. That is why it's important to carefully plan the infrastructure to take into account capacity planning, availability management, and information security aspects.

Workload and Data Size

Workload and size of the data stored and handled must be carefully considered when building your OroCommerce infrastructure.The most obvious parameters to consider for a website are the number of online users and the requests rate. These parameters are true for e-commerce but also affect the calculations for network bandwidth and web-server component capacity. There are also additional factors which impact an e-commerce website like the size of your catalog. Understanding your catalog size will be important when planning your DB server capacity. In addition, one must take into account the number of integrations that will generate Web API loads and the number of transactions that should be handled by application.

Reliability and Scalability

Availability of the e-commerce solution is a crucial business metric because each minute of downtime means users cannot purchase from your website which directly impacts revenue. However, it is practically impossible to have 100% availability. You must take into account that each server can fail and stop processing requests. So when planning your infrastructure, it is important to build a reliable solution using non-reliable components.

Redundancy is the most popular and well-known approach but it’s not cheap. OroCommerce is built using a stateless backend meaning you can easily add 2 and more nodes and configure load balancing using ngnix, haproxy or any other TCP or HTTP balancer including hardware solutions like F5 balancers. Oro recommends using a round robin balancing approach combined with component health checks for all nodes (except the DB).

In addition, the balancer makes it possible to use autoscaling solutions available in popular IaaS platforms like Amazon Web Services (AWS), Microsoft Azure or Google Cloud Platform (GCP).

All 3rd party open-source components used in OroCommerce - PostgreSQL, RabbitMQ, Elasticsearch - support clustering, horizontal scaling and other methods for building both scalable and highly-available solutions. We recommend you consult with corresponding documentations and follow the vendor's best practices.

Monitoring

Monitoring is an important component for operating any website. E-Commerce solutions in particular need monitoring processes that are well-designed and detailed. Monitoring will help everything from detecting and preventing cybersecurity incidents to pinpointing where incidents may have occurred using forensic data and help to identify performance bottlenecks. Be sure to take monitoring needs into account when planning your infrastructure.

We recommend using open-source, industry-recognized monitoring tools like Zabbix or Nagios. Monitoring services are also available in IaaS platforms - like GCP StackDriver, AWS CloudWatch or Azure Monitor. A customer can also use any existing monitoring system including proprietary ones.

While it's important to monitor system parameters like CPU, network utilization, free storage space, etc., it is highly recommended to keep track of the OroCommerce application:

- Check logs for errors and error rates.
- Check availability of the API.
- Calculate page response times using the web server logs.

OroCommerce provides flexibility in configuring any location for logs, logging levels, and more. It also provides the ability to use any existing log forwarding tool - like rsyslog and syslog-ng. Please refer to the OroCommerce installation guidelines for more details on log configuration.

Other OroCommerce application components like DB server, message broker and search engine also need to be monitored. Because of our open-source model for development and distribution, 3rd party tools used by OroCommerce - such as PostgreSQL, RabbitMQ, and Elasticsearch - have many different types of monitoring tools. Supported options can range from standard commands to built-in application monitoring dashboards to ready-to-use plugins for external monitoring systems like Zabbix or Nagios.

Security is a Critical Factor

E-Commerce websites are targets for hackers because they contain a constant flow of transactional payments performed by these sites. By hacking into an e-commerce site, intruders can gain access or harvest sensitive data like Primary Account Numbers (PAN), Card Verification Values (CVV), expiration dates, and other cardholder data. Hacks and data breaches on e-commerce websites can lead to considerable consequences resulting in financial loss for both the site owner and end-customer.

Oro uses best security practices and follows OWASP recommendations for OroCommerce architecture planning, development, and testing processes. Also, standard OroCommerce deployments use payment gateways and do not store any card-holder data (CHD) anywhere in the application, DB or file system.

However, there are still information security best practices to consider when hosting the OroCommerce application:

- Customizations - OroCommerce is the flexible platform that allows any custom features to be added. Keep an eye out for customizations that may store card-holder data. Be sure to place appropriate safeguards for these customizations to prevent CHD leaks or unauthorized modifications.
- PII (Personally Identifiable Information) - Even the standard configuration of OroCommerce stores some information which is classified as PII - like names, addresses, e-mails, phone numbers, etc. PII needs to be carefully handled and stored according to national and international laws and standards.

So even with a state-of-the-art e-commerce solution and no storage of CHD, information security is still an important factor to take into consideration. It’s important to understand that the infrastructure of an application not only defines scalability and reliability of a solution - it also plays into an eCommerce website’s security.

Oro’s cloud solution is PCI DSS compliant. We advise our customers to use this industry standard recommendation even if an OroCommerce implementation does not require PCI DSS compliance. There are number requirements that come from this compliance standard which will help build secure and reliable infrastructures for e-commerce sites.

Security Perimeter

It's highly recommended to only expose needed network nodes and allowed protocols to the public. This will protect websites from unwanted attacks and it will hide implementation details from possible intruders.

A typical deployment recommendation is to use a firewall as an entry point that's exposed to the world with PAT/NAT rules to translate external requests to internal network addresses assigned to the specific servers.

Network Segregation

This technique helps minimize the impact of a possible adversary intrusion to e-commerce sites by isolating it from other applications at a network level. It also places different tiers of an e-commerce application into dedicated subnets. Traffic between subnets go through limited and controlled network interfaces which means that possible intruder cannot get access to other hosts using ssh and other protocols.

For example, a DB server should be placed in a separate subnet with only a DB-connection protocol enabled while also preventing outgoing traffic to other DB servers and other hosts outside of the DB network.

One Role - One Server

One PCI DSS requirement is to use one server for a single primary role. This requirement not only prevents security misconfigurations but also improves manageability of the infrastructure and helps fine tune the performance capacity for each application component.

Performance for different OroCommerce components depends on different server resources:

- Application - depends on CPU, RAM and storage read operations;
- DB - depends on read/write storage operations, RAM, CPU;
- Message queue (RabbitMQ) - depends on RAM;
- Elasticsearch - depends on CPU, RAM and storage read operations.

Also, keeping each component as a dedicated host will help isolate problems when resolving performance issues or troubleshooting incidents.

Sample Deployment Configurations

According to the considerations above, Oro recommends following our standard deployment infrastructure.

Please note that the following configuration diagrams are versions of the OroCloud deployment. Since Oro uses Google Cloud Platform, the diagrams use GCP graphical notations. But, as mentioned earlier, OroCommerce can be deployed on any hosting platform, IaaS, or even bare metal servers.

Minimum Deployment Configurations

For a bare minimum configuration, Oro recommends using at least a dedicated DB server for manageability and performance purposes. This configuration does not satisfy security requirements and you should maybe consider putting it behind the firewall .

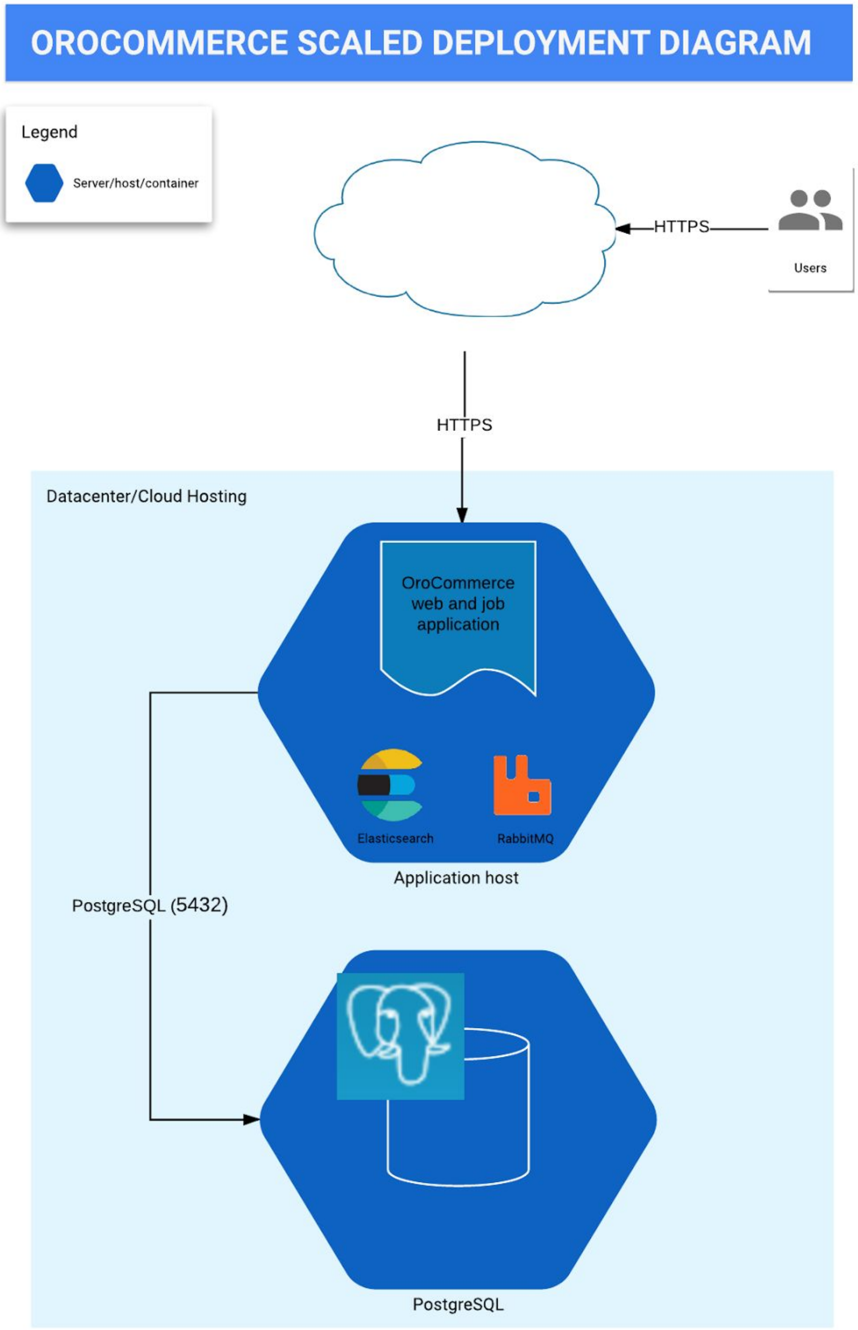

Scalable Deployment Configurations

To get horizontal scaling, you need to move each component of the application to the dedicated node, add redundant nodes and place balancers for traffic distribution between them. This model will also simplify network isolation since the 1st level balancer can also be used as a DMZ or WAF.

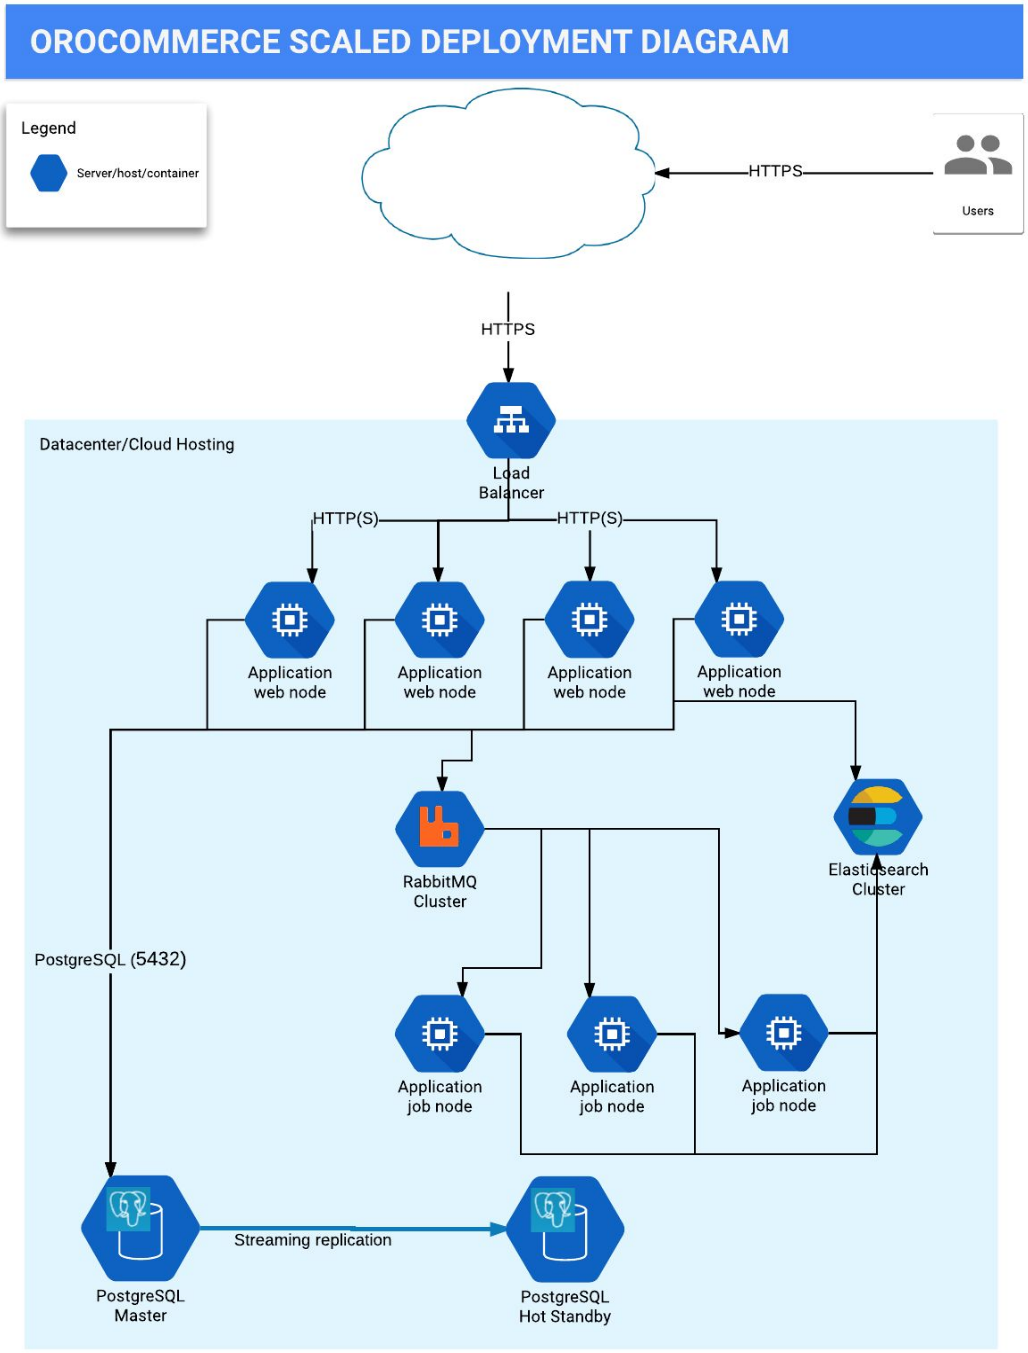

Secure Deployment Configurations

This model strengthens the security of your e-commerce site. In fact, this diagram is exactly like the previous diagram only with every component tier moved to the dedicated subnet.

This PCI DSS compliant approach delivers great manageability options, isolates tiers, mitigates possible adversary intrusion.

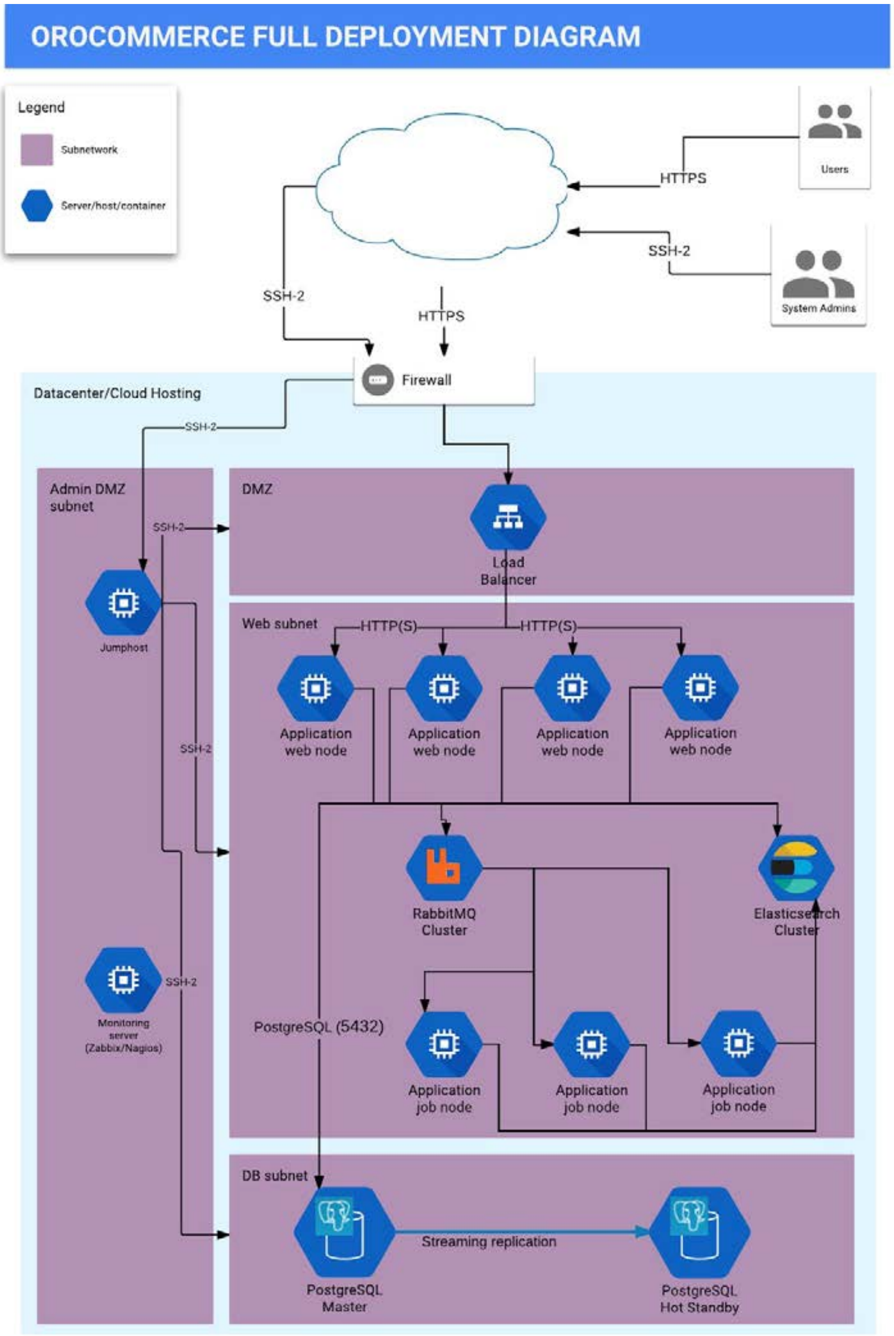

Conclusion

It is important to carefully plan your infrastructure before deploying an e-commerce solution. OroCommerce is the flexible application with a state-of-the-art architecture that can be deployed using any of the provided model. Oro experts can not only help with deployment inside Oro SaaS or PaaS cloud but are always ready to help with building reliable and secure e-commerce site no matter what underlying infrastructure is used - Amazon Web Services (AWS), Microsoft Azure or Google Cloud Platform (GCP), in-house private cloud or even bare-metal servers.