Infrastructure Advice

This document covers the infrastructure recommendations for VoipNow.

Fast Track
If you want to skip reading this, keep in mind that you can choose:

- **Single Node Deployment** - if you do not predict to have more 2,000 users on short term
- **Distributed System Deployment** - if you will have more than 2,000 users soon

Both options can take advantage of High Availability.

Overview

VoipNow is designed to maximize the cloud advantages. The cloud concept is built into the platform from bottom to top.

VoipNow uses raw computing resources such as CPU power, memory, storage space (the so-called IaaS - Infrastructure as a Service) and transforms them into a powerful communication and collaboration service. The communication and collaboration service is designed to be delivered as SaaS (Software as a Service), enabling you to serve:

- private users in small business or large enterprise environments
- the general public (private or business users) in a service provider environment

VoipNow uses the same cloud concepts no matter if you deploy it for your own company or to deliver services to third-parties, in a public or private cloud.

The service levels you would want to enforce may vary, however the principles described here are universal.

Infrastructure

These recommendations take into consideration the type of infrastructure you are going to use.

Your own infrastructure

We consider that you are running your own infrastructure when:

- You have low-level management access to the infrastructure - you have access to server hardware to install virtualization OS.
- The location of the infrastructure is not important - can be on-premises or hosted in a third-party Data Center.
- You pay for raw resources - this means that you are not paying for computing instances on a pay-as-you-go mode.

It does not matter if you have already built a private IaaS cloud with management facilities.

Infrastructure as a service

We consider that you are using a third-party Infrastructure as a Service (IaaS) cloud when:

- All you can see and manage in the infrastructure is the computing instance(s) and storage.
- You lease computing instances (hourly, monthly etc.) and you do not have management access to the hardware they run on.

This is the modern view on the infrastructure services, with many similar popular services like Amazon AWS, Microsoft Azure, Google Cloud, Oracle Cloud etc.

Service distribution

Any private or public cloud service provides CPU power and memory on computing instances. Software requiring more CPU power and memory must be able to use these computing instances efficiently. And this is where service distribution comes in.

Single vs. distributed

VoipNow can be run in two deployment modes:

- on a single computing node
- distributed on multiple computing nodes

Of course, it is entirely your choice on how you run VoipNow. Ideally, this choice should be made at installation time, due to the infrastructure requirements. Even later, assuming that you understand all implications, you will be able to switch from a single computer installation to a distributed one, but under no circumstances will it be possible to switch from a distributed environment to a single computer.

VoipNow features **advanced service distribution** and **cloud management**. Thus, VoipNow is able to use the cloud resources in order to scale infinitely with the number of customers.

Running a distributed system is not trivial, even if the software provides highly integrated management features. It requires advanced knowledge about system administration, security and infrastructure monitoring. The distributed setup is not appropriate for all deployments, therefore we try to guide you in this chapter on how to make the right choice.

High availability

There is a common confusion between service distribution and high availability. The general perception is that, when a system runs on multiple computing nodes, it is also more resilient to failure. In most cases, this is not true. In fact, statistically, service distribution increases the chances of failure.

If you are looking for high availability only, you should not embrace service distribution. Instead, follow the recommendations below.

Your own infrastructure

Physical level

High availability must be built starting with the physical environment. Your minimum requirements should be:

- redundant power supply
- redundant L2 network connections
- redundant L3 routing on the premises
- multi-homed Internet connections
- infrastructure monitoring and incident detection
- reliable server hardware with redundant components
- redundant storage system

It is expensive to provision such an environment on your own, this is one of the reasons why it is recommended to run your infrastructure in a professionally managed Data Center with a cloud service provider.

Software level

The recommended approach is to guarantee high availability by deploying HA at the virtualization OS level. Each vendor provides its own solution:

- VMware vSphere - VMware vSphere High Availability (HA)   **BEST**
- Citrix XenServer - High Availability for Citrix XenServer
- Virtuozzo - RHCS Cluster

We do not recommend virtual machine OS level high availability due to the management overhead this configuration imposes, especially on migrations. While it is possible to provision VoipNow this way, it is not recommended.

Infrastructure as a service

Many service providers have special computing instances that ensure a better availability. In their infrastructure, these are provisioned using virtualization system level high availability, which is already a recommended way to run a highly available VoipNow system.

Capacity scaling

The only reason to run a distributed system is **because you need massive scaling capabilities and you still want to manage a single system.** If you have to scale and you do not mind managing multiple systems, it is possible to provision multiple VoipNow systems with each of them running on a single computing node independently.

Before anything else, it is necessary to understand the limits of the software running on a single computing node.

The resources used by the system scale with the number of users and the features used. Therefore, it is impossible to give exact figures for the VoipNow system performance on a single node. However, generally speaking, a modern computer can provide services for more than 2,000 users. Public clouds provide such instances (e.g. Amazon Extra Large instance).

When VoipNow is employed in applications that use intensively telephony functions, more than 1,000 concurrent calls can be provisioned.

When Distributed?

If you think that your system will have to handle more than this on a short term and you still want a single system, it is recommended to start with a distributed setup. Otherwise, a single node deployment will reduce your operational costs.

Storage

While database layers require a fast local storage, it's better to rely on cloud storage for storing file resources like images, documents, recordings, voicemails, faxes etc.

VoipNow implements an easy to setup Amazon S3 compatible client interface. It makes files storage very economical, scalable and fast.

What's next?

Based on your infrastructure of choice, the next step is to understand the requirements:

- Single Node Deployment
- Distributed System Deployment

We hope that you have already decided what infrastructure is most suitable for you.

Get Help
If you still have trouble in making that decision, do not hesitate to ask for advice from our Professional Services team and follow VoipNow administrators in our community.

If you need more help, our expert engineers are waiting for you in the Support Zone.