Project Wendelin

The Wendelin project started in the beginning of 2015 with Nexedi as consortium leader in charge of managing the development of a Big Data Solution "Made in France". Wendelin will be based on open source software components and combines a number of widely used libraries for data analysis, modification and visualization. The project also includes the development of prototype applications in the automotive and green energy sectors underlining it's purpose of being immideately applicable for the development of industrial solutions.

Meet The Wendelin Stack

The Wendelin stack is written in 100% Python. It leverages SlapOS for cloud deployment and NEO for distributed storage while ERP5 is used as platform to interconnect the various libraries available, enable the creation of web-based visualization applications and allow to extend Wendelin more towards business processes ("Convergence Ready"). One of the core features of Wendelin is it's out-of-core computation capability, which will allow Wendelin based stacks to easily extend computation capacity beyond the limits of available hardware in a cluster. Of equal importance is the interface with Scikit-Learn providing core machine learning capabilities to all Wendelin-based applications.

New Features in 0.4 alpha

The new release of version 0.4 alpha includes a lot of "under-the-hood" features making working with Wendelin much easier. A lot of effort has gone into getting the Wendelin stack to install faster. With this release, the installation time on Debian 8.1 (64bit) has been reduced from 4 hours to around 30 minutes. We also switched the default installation routine from a single zope node setup to using a cluster of nodes for Wendelin and now provide a fully functional development instance able to run live tests and allowing anyone to develop on top of Wendelin. On the technical side, our "Wendelin out-of-core" functionality has been updated to version 0.4 after fixing a ZODB invalidation bug occuring in heavy loaded cluster environments. Finally, some new tutorials have been added to the examples section showing how to easily get started with Wendelin.

What's Next?

With version 0.4 out of the way, Nexedi is already knee-deep into development of the next release. We have successfully integrated Jupyter's IPython Notebook into Wendelin and are now working on integrating it as a configurable feature in our official release. Work has also begun on adding Pandas for more visualization options, athough we are not sure the latter will already make it into the next release. Lastly, we are also looking into fixing some ZODB size related issues, ,so there are a lot of nice things and improvements in the immideate pipeline.

Tutorials: Getting Started With Wendelin

Wendelin is still under heavy development, but at this point it is already possible to get a working instance and start playing with it. The following steps are still bound to change as Wendelin matures, but if you want to give it a try, read along or follow the detailed instructions on how to get started with Wendelin and how to configure your Wendelin instance.

1. You will need a machine with at least **4GB RAM**, **20GB disk space** and a **Virtual Machine** installed - preferrably Debian 8.1 (64bit).
2. You can follow the instructions provided for VMWare and Virtualbox to setup your VM on Debian.
3. Once you have your VM, run the following (root permission required):

   ```
   root@debian8:~# wget http://deploy.nexedi.cn/wendelin-standalone
   root@debian8:~# bash wendelin-standalone
   root@debian8:~# chown -R slapsoft:slapsoft /opt/slapgrid
   ```

4. You can monitor your build progress by either one of the following:

   ```
   root@debian8:~# watch -n 30 erp5-show -s
   root@debian8:~# tail -f /opt/slapos/log/*.log
   ```

5. To check if your instance is ready, you can:

   ```
   root@debian8:~#  erp5-show -s
   ```

which, once done, should return:

```
Build successful, connect to: https://zope:insecure@[2001::bd4d]:16001
```

6. This IPv6 is an internal one for your machine and the only way to access it is to run the browser inside the VM machine. To access the Wendelin instance from outside the VM machine it is currently still necessary to do the following:
7. Find the internal IPv4 on which Wendelin is listening to:

```
root@debian8:~# ps xa | grep runzope
root@debian8:~# vi /srv/slapgrid/slappart0/etc/zope.conf
(and fine block with

  address 10.0.210.201:12001
)
```

8. Create a ssh tunnel like above from host linux machine:

```
Host debian8_tun
HostName debian8
User root
LocalForward 2200    10.0.210.201:12001
```

9. Access locally http://localhost:2200/ and use for username: **zope** and for password: **insecure**
10. There are some caveats when installing Wendelin outlined in the linked tutorial. We are trying to adress those as development moves along. Make sure to also check the configuration page to see how to get started once you are up and running.

Summary

Wendelin has just been updated to version 0.4 alpha bringing some performance improvements, bug fixes and new libraries being available. We showed you how to get your Wendelin instances along with some tutorials to configure your stack. Stay tuned to this blog as Wendelin evolves and hopefully becomes the go-to open source platform for working with Big Data.