Secure Election Infrastructures Based on IPv6 Clouds
- *Booklet* -

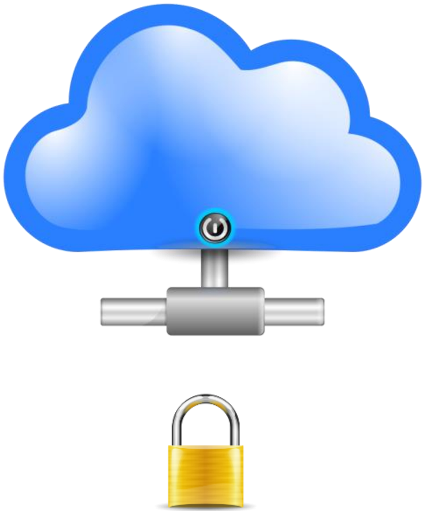

GOVERNMENTS ENABLED WITH IPv6
GEN6

ICT PSP
ICT POLICY SUPPORT PROGRAMME
part of the Competitiveness and Innovation Framework Programme CIP

cip competitiveness and innovation framework programme 2007–2013

Author: Gabriela Gheorghe (UL)

Disclaimer

The GEN6 project (number 261584) is co-funded by the European Commission under the ICT Policy Support Programme (PSP) as part of the Competitiveness and Innovation framework Programme (CIP). This document contains material that is the copyright of certain GEN6 partners and the EC, and that may be shared, reproduced or copied "as is", following the Creative Commons "Attribution-NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-NC 3.0) licence. Consequently, you are free to share (copy, distribute, transmit) this work, but you need to respect the attribution (respecting the project and authors names, organizations, logos and including the project web site URL "http://www.gen6-project.eu"), and use the document for non-commercial purposes only, and without any alteration, transformation or building derivatives upon this work.

The information herein does not necessarily express the opinion of the EC. The EC and the document authors are not responsible for any use that might be made of data appearing herein and effects that result from doing so. The GEN6 partners do not warrant that the information contained herein is capable of use, or that use of the information is free from risk, and so do not accept liability for any direct nor indirect loss or damage suffered by any person using this information.

Cloud and IPv6 make sense together

For public administrations engaged in building future-proof infrastructures, IPv6 and cloud systems managed in-house are worth to consider together.

Both cloud infrastructures and IPv6 are technologies of the present and it is time to make them part of the future. Why? IPv6 itself enables addressability and thus helps end-to-end connectivity, when it comes to many heterogeneous computing resources (for example, mobile devices and Internet of Things). Cloud systems help reduce spending on infrastructure, improve accessibility, and enable scaling. Cloud software to deploy and manage fleets of virtual resources is already available, either proprietary or open-source; this technology is already offering management features that network administrators where only dreaming about before. Together, cloud and IPv6 make sense together because the resources that IPv6 can access, can be virtualised in the cloud and controlled remotely.

Election pilot based on an IPv6-enabled cloud

The University of Luxembourg (UL) hosted the first IPv6-only cloud in a production environment, in cooperation with Citkomm and Nephos6. This cloud system has served the May 2014 elections in Germany for the Citkomm customers, and successfully served 5% of the requests of all the citizens accessing the election results presented there.

The service showcased by this pilot is the presentation of the election website. Throughout the election days, citizens of various municipalities in North-Rhine Westfalia could access the current voting count on a Citkomm-hosted website (http://wahlen.citkomm.de/). The backend webserver for this website has traditionally been IPv4-only, and with this pilot we showcased two novelties at the same time:

- IPv6 enablement of website needed especially by those citizens accessing it from machines that are IPv6 enabled;
- Cloud computing assurance (availability, resilience, scalability, security) when it comes to handling large amounts of user traffic.

The pilot achieved its purpose fully. It employed an open-source cloud distribution, OpenStack Havana, that was adapted in-house to support IPv6, as native IPv6 support is not yet official. The pilot has passed an intensive testing phase that covered heavy load generation and handling. During the test and production phase, the pilot was subject to extensive QoS monitoring and performance data collection.

The pilot gathered together a number of presentation servers on the Citkomm site, and two others on the Luxembourg site. This is shown in the figures below. The first figure shows how resources from different locations (one is QSC, the other is the Luxembourg site) were integrated together under the same *citkomm.de* domain. The QSC resources are all IPv4-only, while the Luxembourg ones were addressable only via IPv6.

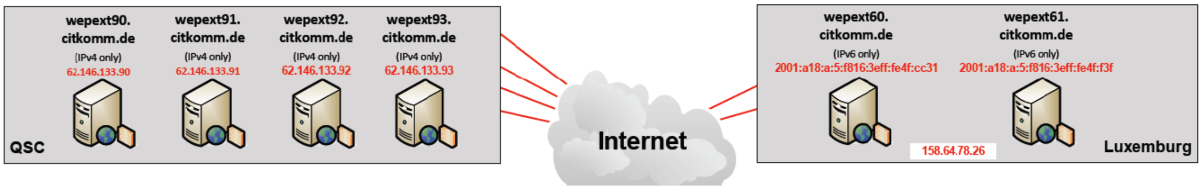

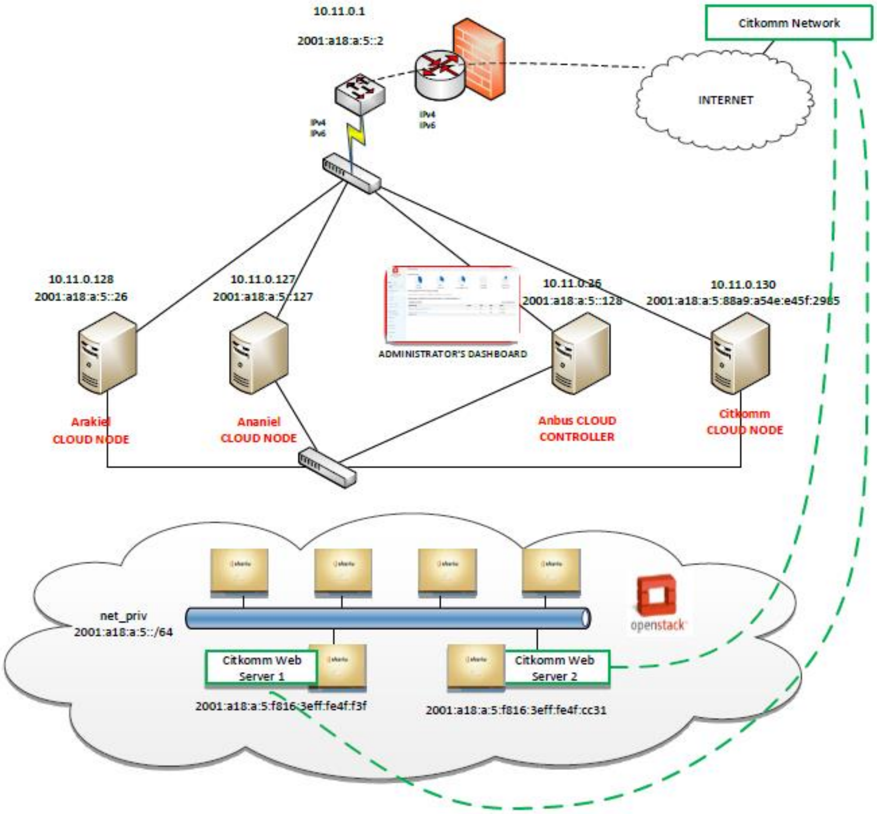

The second figure above shows the deployment at the UL site in more details: on 4 physical servers (of which one is a cloud controller and the other 3 are cloud nodes) there are a number of virtual resources "in the cloud". These resources are virtual machines: images of complete operating system and applications running on top of it. These virtual resources are managed internally by the cloud operating system – OpenStack in this case – and are situated in the same network segment in the university network, protected by a firewall.

Transition to IPv6

The transition to IPv6 can be of the application to be virtualized and deployed in the cloud, of OpenStack system itself.

In our case, the application to be run in the cloud was the presentation website. This was HTML code and hence was independent of the IP protocol understood by the browser.

The integration of IPv6 in OpenStack is not yet officially achieved in the open-source community. At UL, the OpenStack Havana testbed has been patched for full IPv6 support with the help from Nephos6, an IT company based in Raleigh, USA. All details of the patch can be found in a previously published whitepaper (http://www.nephos6.com/pdf/OpenStack-Havana-on-IPv6.pdf) and they cover the address assignment, and some routing issues in OpenStack. The patch is relatively easy to deploy and, once installed, it is possible to launch virtual machines with native IPv6 addresses. The patch will be officially integrated in the next version of the OpenStack software.

When virtual machines can have IPv6 addresses, they can be accessible directly, without any need for intermediate (proxy) configurations, by both users and network administrators. In other words, everything that is set up within these virtual machines becomes immediately accessible to everybody. Think of a virtual machine as a virtual computer, where any application can run to serve user, and communicate with other virtual machines to achieve a common purpose.

Complex infrastructure monitoring at your fingertips

As typical cloud-based software goes, OpenStack gives a very granular way to manage virtualized resources. Virtual resources in this scenario host the webserver of the election results website, and here they are virtual machines with a Linux operating system on top. Some of the management features offered by OpenStack for virtual machines cover:

- Virtual machines can be switched on, off, can be paused, can be replicated at various states in their lifecycle,
- Virtual machines can be firewalled in different ways from the OpenStack dashboard,
- Virtual machines can change network configurations (one-by-one or in groups). For example, machine instances can be assigned different virtual IPv4 or IPv6 addresses,
- Virtual machines can be monitored individually or in groups at hypervisor level,
- Virtual machines can be made to execute script actions at bootup,
- Virtual machines can be migrated from one physical host to another, without losing state.

Below you can see one view from the OpenStack dashboard – the “control room” of the cloud – from where the cloud administrator can visualize the existing resources, tune them, or change various parameters of the infrastructure. These features are included in the out-of-the-box OpenStack software.

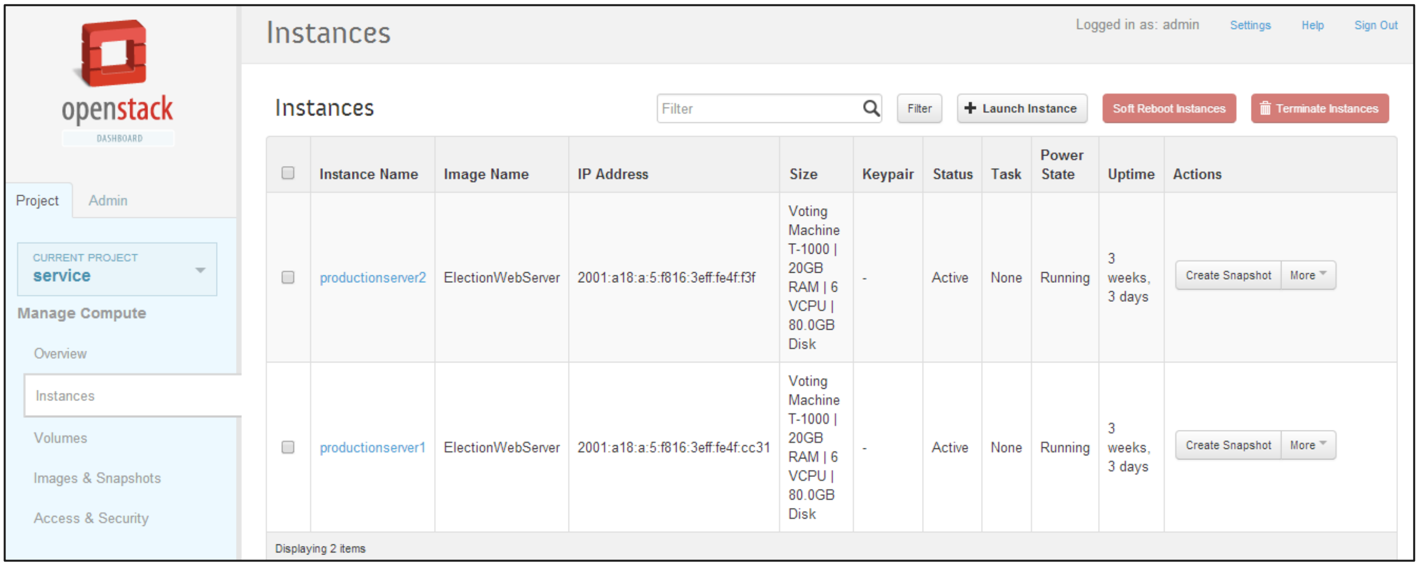

Instances

□	Instance Name	Image Name	IP Address	Size	Keypair	Status	Task	Power State	Uptime	Actions
□	productionserver2	ElectionWebServer	2001:a18:a:5:f816:3eff:fe4f:f3f	Voting Machine T-1000 | 20GB RAM | 6 VCPU | 80.0GB Disk	-	Active	None	Running	3 weeks, 3 days	Create Snapshot More
□	productionserver1	ElectionWebServer	2001:a18:a:5:f816:3eff:fe4f:cc31	Voting Machine T-1000 | 20GB RAM | 6 VCPU | 80.0GB Disk	-	Active	None	Running	3 weeks, 3 days	Create Snapshot More

Displaying 2 items

Therefore, with cloud features, e-government infrastructures can be managed better than ever before:

- the administrator can access virtual resources, bare-metal system and network information at any time and at different levels of granularity (virtual machines, virtualization level, hardware monitors on the physical machines on which the virtual resources reside),
- scaling (up or down) of resources can be achieved at runtime with the press of a button, since elasticity is one of the main features of cloud systems, and is implemented in OpenStack as is,
- essential information in the area of runtime QoS monitoring and assurance.

As a proof of concept, UL has experimented on the monitoring features that already exist and that can be added on top of this cloud distribution. This work was done together with Citkomm and Nephos6. We have instrumented the typical OpenStack monitoring so that an e-government infrastructure administrator can be offered more information about the distributed system in a centralized way. In this work, and in the following figures, we have used Sonar, a Nephos6 Service Assurance tool.

During the actual elections, we monitored the end-user experience from several locations in Europe and in the US. We measured HTTP response time throughout the election day for all virtual resources that users could access (one single URL could direct end-users to the UL infrastructure when the connection was over IPv6), and correlated this time with some information from within the UL deployment. This approach is useful for several main reasons: first, the administrator can have a concrete idea of the user-side experience of the application running on top of the virtualized infrastructure. This can be seen in the figure below, showing the user-side experience of the Citkomm URL http://wahlen.citkomm.de/ during the election on the 25th of May, between 8pm and 9:45pm.

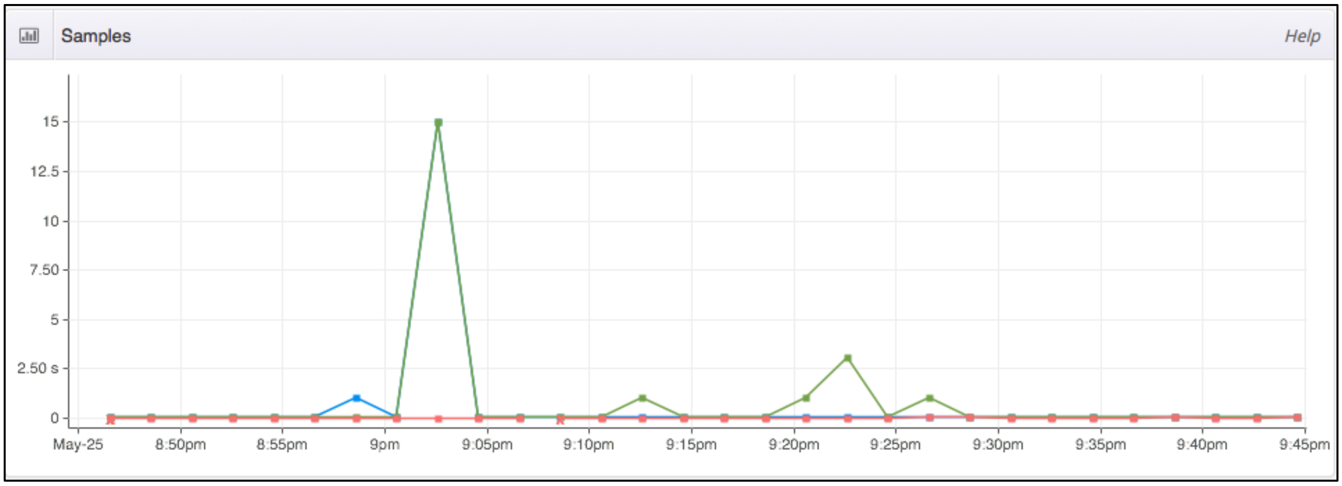

Second, the administration can compare user experience from different locations. The figure below shows how, in our proof of concept, the administrator can experience a cloud dashboard: a map of Europe with the marks for where the monitoring scripts are deployed, and a diagram showing the performance experienced from the different locations when accessing the resource in the cloud (in this case, the elections website at the URL indicated above). This information can provide hints about potential problems (e.g., it is likely a network problem if the user experience is bad from some locations, while it can be a server-side problem if the user experience is bad from all locations). For example, in the figure below the red line, corresponding to the Gen6 OpenStack locally-deployed measurement script, has a much better performance than the other two, which are subject to network delays associated with their location.

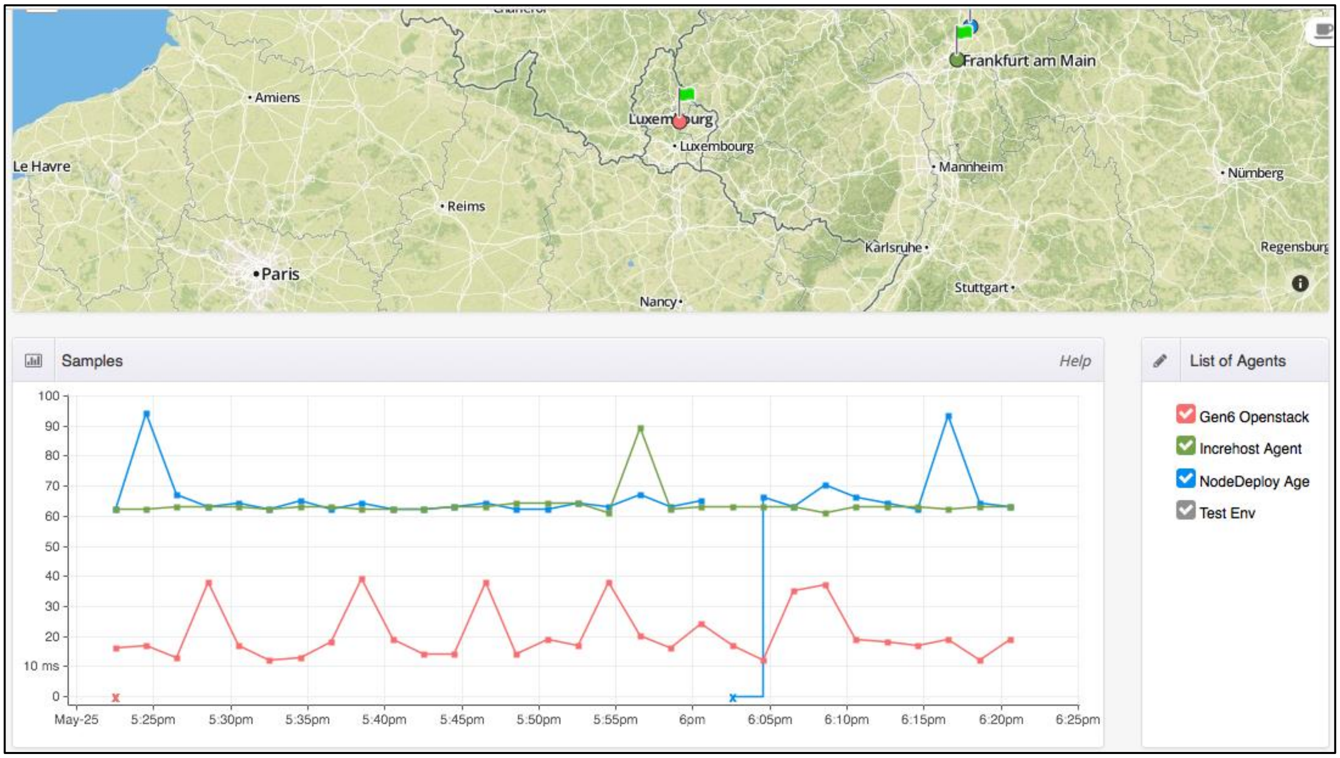

Third, the administrator can *act* on the observed issues, whether by investigating at the server-side, or by moving resources from one virtual/physical network to another, or by starting up new resources. These possibilities are not available in traditional networks.

In our proof of concept, we could visualize different types of monitoring data in the same dashboard:

- HTTP data and ping, for the application running on top of the cloud infrastructure,

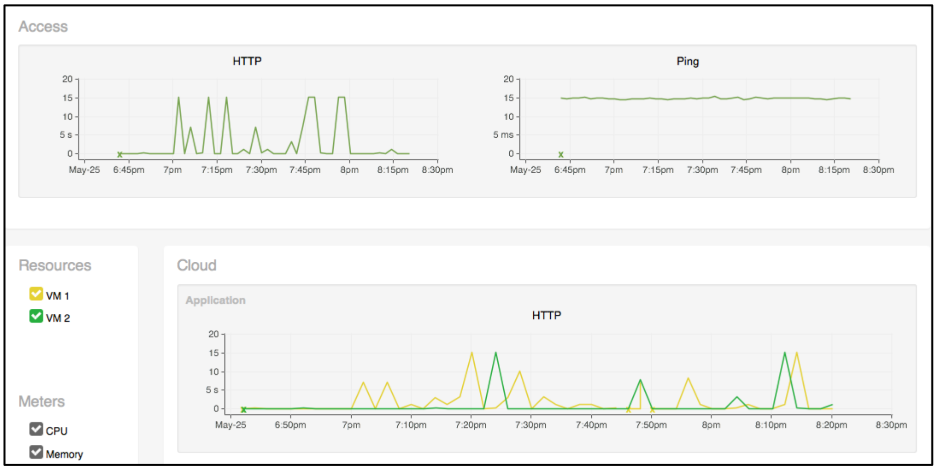

- Data reported by monitoring tools such as Munin, that look at disk, CPU and network operations when virtual machines run,

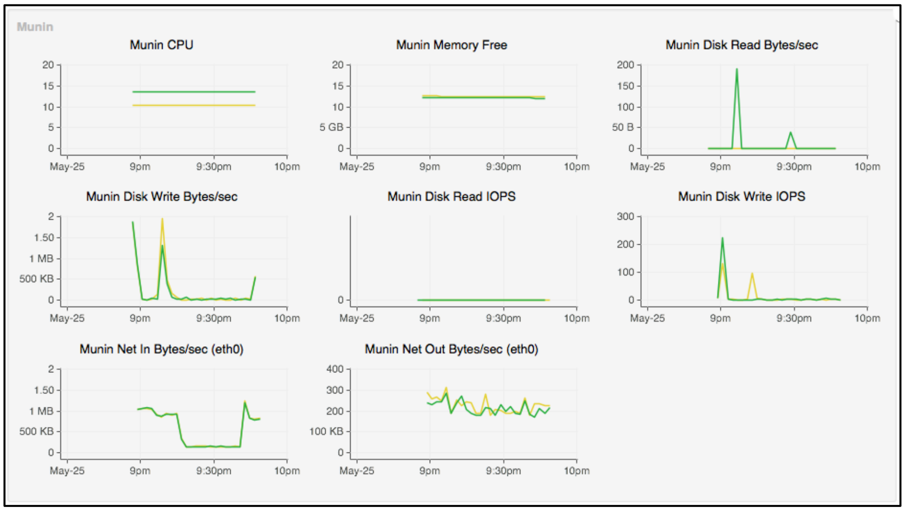

- Physical infrastructure data (the physical machines on which the virtual resources are running) that is gathered and reported by a service within OpenStack called Ceilometer, which is in charge of reporting system statistics.

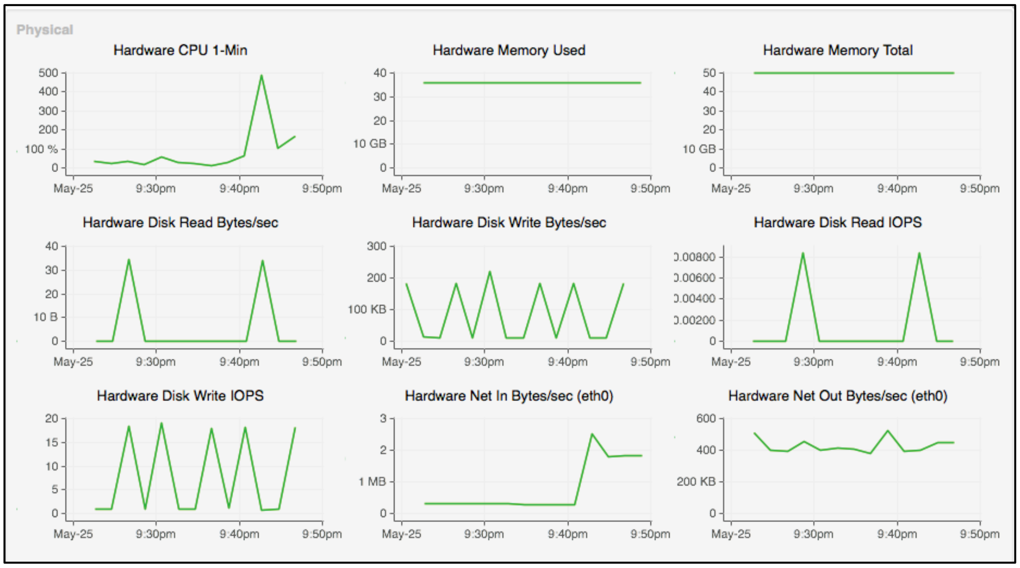

This rich palette of data is extremely useful for the infrastructure administrator, as it can be used for root-case analysis and mitigation in case there is an infrastructure incident. Moreover, with multiple virtual resources addressable by IPv6, the troubleshooting process is becoming easier because addressing is now straightforward: every virtual machine can be directly accessed and queried. With traditional IP addresses, the administrator had a much harder time to configure NAT and firewalls on individual middleboxes in the infrastructure.

Security considerations

Security requirements of the election infrastructure in our experiment cover several aspects:

- Firewalling over the physical and virtual resources, to prevent unauthorized access at network level. Even though IPv6 eliminates NATs, the need for a correctly configured firewall is just as dire when using IPv6 as when using IPv4,

- Access controls when it comes to accessing the virtual resources (e.g., the website and its updates, the backend connection, log data backups),
- Resource exhaustion detection / prevention are particularly needed in a cloud context in which illegitimately demanding resources can exhaust the physical capabilities, because the cloud's inner elasticity mechanisms would be triggered easily,
- Security monitoring and testing, which are parts of common security and reliability practices.

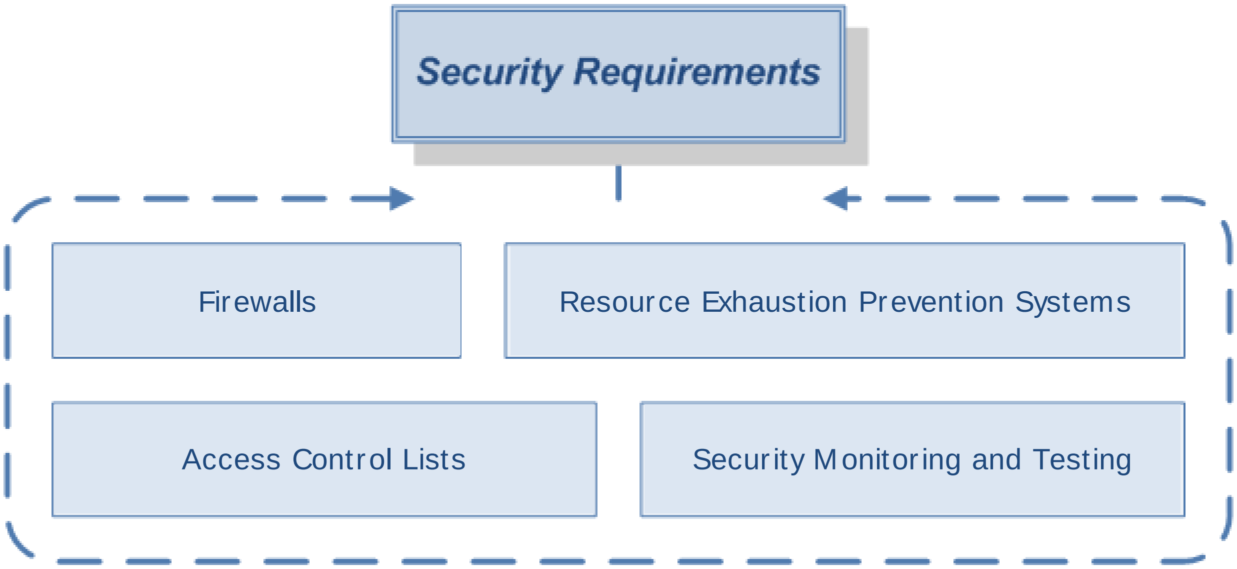
Security Requirements
Firewalls
Resource Exhaustion Prevention Systems
Access Control Lists
Security Monitoring and Testing

Firewalling in OpenStack can be done from the main administrator dashboard, the "Security groups" settings tab. The figure below shows how an administrator can view and edit security group rules for virtual machines and set constraints on the inbound and outbound traffic on IP level and above.

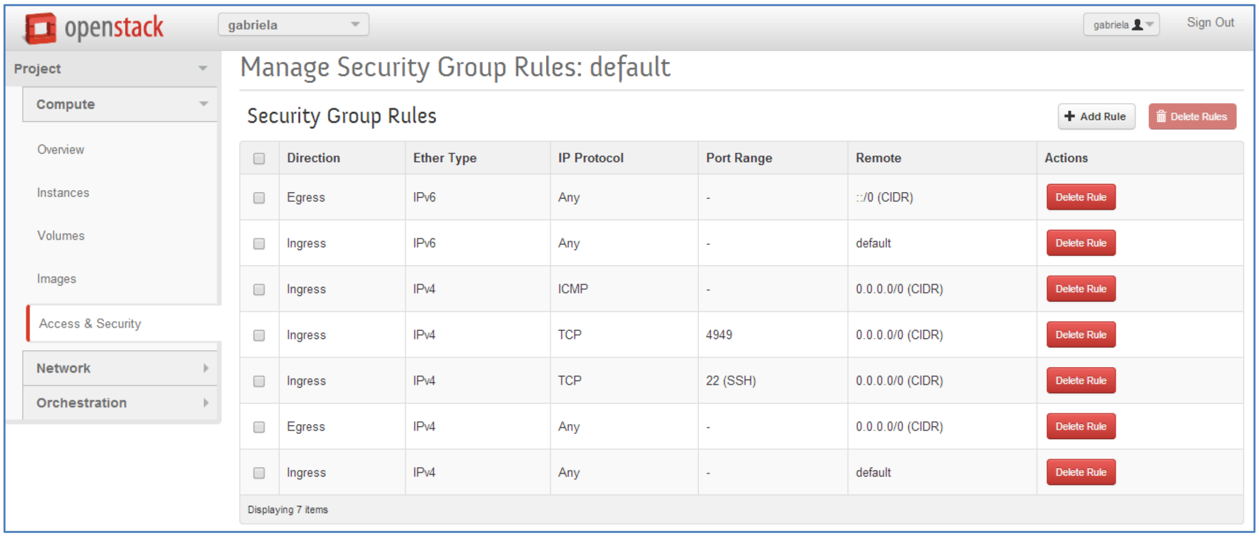

Manage Security Group Rules: default

Security Group Rules

□	Direction	Ether Type	IP Protocol	Port Range	Remote	Actions
□	Egress	IPv6	Any	-	::/0 (CIDR)	Delete Rule
□	Ingress	IPv6	Any	-	default	Delete Rule
□	Ingress	IPv4	ICMP	-	0.0.0.0/0 (CIDR)	Delete Rule
□	Ingress	IPv4	TCP	4949	0.0.0.0/0 (CIDR)	Delete Rule
□	Ingress	IPv4	TCP	22 (SSH)	0.0.0.0/0 (CIDR)	Delete Rule
□	Egress	IPv4	Any	-	0.0.0.0/0 (CIDR)	Delete Rule
□	Ingress	IPv4	Any	-	default	Delete Rule

Displaying 7 items

Advanced access controls for authentication and authorization are available in OpenStack's command line interface. OpenStack's identity service can be connected with an LDAP server, external multi-factor authentication services or Kerberos systems. A super-administrator can create accounts and associate permissions to what OpenStack calls *tenants* – isolated projects (i.e., sets of virtual resources managed by a single administrator) in the cloud. Tenants are subject to quota controls (e.g., number of virtual machines they can launch, number of processor cores they can occupy, IP address space they can use, disk space, etc). Tenants restrict, therefore, a user’s access to certain virtual resources; access key pairs for resources are available per user, but, as the OpenStack manual mentioned, quotas control resource consumption of each tenant across hardware resources, to ensure tenant isolation. Advanced logging and monitoring features to view user activity are also available.

For cloud tenants such as election testbeds, resource exhaustion events can be highly damaging because they affect the election website and hence citizens will no longer be able to access it. The tenant-based design that comes natively with OpenStack can isolate damage from one tenant to another, hence the spreading of the problem is limited to the physical resources that the particular tenant is using. Nevertheless, exhausting resources within a tenant stays problematic, and that is usually brought by Denial of Service (DoS) attacks. In the election scenario, UL has been considering how analytics on monitoring data can be used to enable reliability and system security in the face of DoS attacks. In our approach, by periodically probing the election website it can be inferred if the website is accessible from all virtual sites; if that is not working as expected (e.g., response time within a given time threshold), it is possible to infer, by distributed monitoring, what resources are underperforming. There are several ways to react to this situation: spawn new resources on the fly and reroute traffic there, migrate virtual machines to different physical hosts, restart virtual machines. OpenStack makes it easy and painless to perform such reactions, depending on the situation at hand.

Conclusions from this GEN6 pilot

With this experiment, we have shown that existing e-government services can be enabled with IPv6 and that open-source cloud distributions can successfully face real-world requirements for the public sector. Moreover, our proof of concept shows that it is possible to integrate cloud-based services into a real infrastructure and add to its scalability, and with OpenStack those operations are now a reality. The resulting system, with a mix of physical and virtual resources working together, can successfully handle real-world peak load, and both IPv4 and IPv6 "islands" can co-exist in the same infrastructure and bring added value. In all, we have shown that it is possible to build future proof infrastructures with both IPv6 and cloud technologies.

Further resources

OpenStack Administrator’s Guide, http://docs.openstack.org/admin-guide-cloud/content/ch_preface.html

OpenStack Security Guide, http://docs.openstack.org/security-guide/security-guide.pdf

Press release on Citkomm-UL-Nephos6 election cloud testbed,
http://www.gen6-project.eu/fileadmin/GEN6/Flyer_GEN6/Pressemitteilung_GEN6_02.06.14.pdf

Citkomm video-report on the May 25th 2014 election (in German),
http://www.citkomm.de/ueber-uns/news/detailansicht/article/video-vom-citkomm-wahlabend.html