Helix Nebula – The Science Cloud

Title: D5.2: Report on future technical requirements

Editor: Phil Evans

Work Package: WP5

Submission Date: 21 March 2014

Distribution: Public

Nature: Report

Revision History:

0.1 Initial Revision

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

Executive Summary

Throughout the Helix Nebula initiative a number of activities have taken place to identify the requirements of the users for a brokerage platform that would meet their particular flagship application needs. Initially these requirements were formulated by the Technical Architecture Committee as a series of User Stories that were created together with the existing user communities of the initiative.

The requirements were also further elaborated and formalised through the activities performed as part of the Helix Nebula FP7 project with the user requirements gathered by Work Package 3.

Through the various Proof of Concept, Pilot and Production deployment activities that have been undertaken of the last two years, users have had the opportunity to assess the suitability of the platform against their requirements. All three flagship deployments have been successfully completed however from the feedback received it is clear that there is still some development of the platform required in order to further align the services offered with the requirements of the users.

A critical analysis presented in this document reviews the User Stories, User Requirements and Flagship deployment feedback received in order to identify a series of User Stories that should be addressed by the technical roadmap of Helix Nebula going forward.

The User Stories identified can be summarised as follows:

- Provisioning: Users can successfully provision virtual resources via both the web-console and API, however further iteration is required in the form of support for a Storage and Disk Management
- Security: Currently there is no support for security configuration via the Blue Box and a series of requirements have been identified going forward
- Image Marketplace: A basic service catalogue has been introduced within the Blue Box but requirements have been identified to support better searching and selection of the services available within Helix Nebula
- Data Management: Data Management requirements have to date not been addressed but are contained on the roadmap for Helix Nebula
- Data Marketplace: The Data Marketplace requirements have not yet been addressed however a H2020 proposal has been submitted called XzelCloud that will address these particular needs

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

- Billing & Service Level Reporting: Evolution of the Billing and Metering functionality will be required in order to address the needs of the users and in addition programmatic access to Service Level Reporting is required.

Going forward, the requirements identified in this document must be further elaborated together with the user communities and suppliers, prioritised, and integrated with the roadmap of the service.

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

Contents

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

1. Introduction

This document performs an analysis against the original requirements for Helix Nebula as defined through groups such as the Technical Architecture Committee and deliverables D3.3 and D5.3 compiled as part of the Helix Nebula FP7 project and a critical review as to the level to which these requirements have been implemented within the Helix Nebula Marketplace.

Furthermore, the document providers a set of new requirements which should become the focus going forward for the Helix Nebula Marketplace that would help further develop and innovate on the services which are offered in the Marketplace whilst furthering the level to which the requirements identified throughout the initiative are implemented.

The document is structured as follows:

- Section 2: An outline of the functionality that should be supported by the brokerage functionality of Helix Nebula.
- Section 3: A record of the User Stories that were identified by the Technical Architecture Committee of Helix Nebula including an analysis of which User Stories have been addressed with the first release of Helix Nebula.
- Section 4: A high level analysis of the requirements that were gathered as part of Work Package 3, an analysis of what functionality has been implemented for each requirements group and specification of a series of User Stories to cover functionality that remains to be implemented.
- Section 5: Additional requirements that have not yet been covered by the previous sections and have been identified as still needing to be addressed.

2. Brokerage Platform Features

Building upon the requirements that were identified by the Helix Nebula Tech Arch committee and documented in the following publically available document (http://www.helix-nebula.eu/index.php/uploads/file/81/33/HelixNebulaArchitecture.pdf.html) delivery D4.3 Cloud Provisioning Report details the features of the Helix Nebula Brokerage Platform (the Blue Box) as having:

1. A single common interface for the customer to interact with all suppliers. From a customer perspective the Blue Box provides a single, common way to interface with all suppliers. The customer can interact with all suppliers through the Blue Box without noticing specific differences in the interface of each of the suppliers. The Blue Box achieves this by providing a common API layer on the customer side of the Blue Box, unlocking all functionality of the suppliers behind it.
2. The user can create configurations of infrastructure environments. The Blue Box should support configuration of infrastructure environments that the user can request from the suppliers at a given point in time. The configuration includes the description of the components and interconnections between the components to form a coherent environment. The configuration of a component describes the (preferred) hardware, the virtual machine image (OS and additional software) and the (preferred) supplier. The configuration of the interconnections describes for each component to which other component it is connected.
3. The provisioning of infrastructure on request of the user. The Blue Box should be able to accept a request from the user, based on a user defined configuration, and interact with one or more suppliers so the suppliers can provision the requested infrastructure to the user. During the creation of the (automatic) request the user can (auto) select the supplier, hardware and image directly from lists of available options in the Blue Box or use the preferences as described in the configuration. The Blue Box sends the request to the appropriate suppliers, tracks the provisioning of the requested components, provides the proper information to create a coherent infrastructure environment and reports the outcome of the provisioning (success or failure and configuration details) to the user.
4. The support of an Image library. The Blue Box should maintain an Image library where users can create and store virtual machine images that can be used by all suppliers during provisioning of infrastructure. The images should be able to contain the operating system, user applications and scripts. It should be able to execute the scripts during provisioning.
5. The support of Identity management and Security towards the suppliers. The user should be able to enter identity credentials (e.g. username and password) in the Blue Box. The Blue Box uses those credentials in the communication with the suppliers during the provisioning of infrastructure. The identity credentials are issued by each of the suppliers, given to each

user that is authorised and stored in the Blue Box to request the provisioning of infrastructure at the issuing supplier. At this point in time there is no requirement for a single identity credential per user that is accepted by all suppliers for the provisioning of infrastructure. It is sufficient when the Blue Box allows entering an identity credential per supplier.

6. The termination of infrastructure on request of the user. The user should be able to terminate (part of) the infrastructure that was previously provisioned to him via the Blue Box. The Blue Box should present all infrastructures currently in use by the user and which was provisioned via the Blue Box. The user should be able to select (part of) the infrastructure for termination and send the termination request to the Blue Box. The Blue Box sends the termination request to the appropriate suppliers, and tracks the termination of the requested components and reports the outcome of the termination (success or failure) to the user.

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

3. Technical User Stories Analysis

This section details a critical review of the User Cases that were initially identified by the Tech Arch group and provides an analysis to the level to which each use case has been implemented for the currently chosen Blue Box of SlipStream.

3.1 Simple Provisioning

Req Id	Requirement	Implementation Level
1.1	As a HN consumer, I can provision a single VM, in a single API call, specifying the reference image to start from as well as resources to be allocated (e.g. CPU, RAM), such that I can access the created VM in the cloud of my choice.	Implemented
1.2	As a HN consumer, I can rapidly provision N VMs (e.g. via thin provisioning / copy-on-write), in a single API call, specifying the same as the user story #1, such that I can access the N created VMs in the cloud of my choice.	Implemented
1.3	As a HN consumer, I can create, in a single API call, a new empty disk/volume, such that I can use it to build a new machine or as an additional disk/volume to a running VM (without the need of rebooting it).	Not Implemented
1.4	As a HN consumer, I can provision as for user story #1, with extra disks/volumes, such that I can keep my reference VM small and change available disk space using extra disks/volumes. The extra disks can be: persisted or volatile.	Implemented
1.5	As a HN consumer, I can create, in a single API call, backup snapshots of my disks/volumes, such that I can re-provision and recover destroyed volumes.	Not Implemented at a Blue Box level
1.6	As a HN consumer I can tag instances & Images, Volumes and snapshots, such that I can easily catalog and indicate their contents.	Implemented
1.7	As a HN consumer I can assign a public and private IP to an instance such that I may boot instances with pre-determined network configurations.	Implemented
1.8	Public IP addresses with reversible DNS.	Cloud Provider Specific
1.9	Compliance with EC2 API.	Implemented - EC2 Bridge availability within the HNX platform.

3.2 Contextualisation

Req Id	Requirement	Implementation State
2.1	As a HN consumer, I can specify when provisioning machines, a public SSH key, such that it is available from within the machine during boot for a script to populate the .ssh/authorized_keys file. This means that I do not require to know username/password of the images I provision in order to access them in a secure way.	Implemented
2.2	As a HN consumer, I can specify when provisioning machines, user data (in string/text format), such that it is available from within the machine during boot for a script to act on this information.	Implemented
2.3	As a HN consumer, I can provision middleware and applications on VMs, using pre-configured templates.	Implemented

3.3 Image Management

Req Id	Requirement	Implementation State
3.1	As a HN consumer, I can list public images, or images belonging to me, such that I can discover what images are available and select the right image to instantiate.	Implemented
3.2	As a HN consumer, I can delete images belonging to me, such that I can keep my list of images clean and tidy.	Implemented
3.3	As a HN consumer, I can convert an existing image to the cloud specific format and contextualisation, following clear instructions and available tools, such that I can do this on my own, without requiring specialist support from the cloud provider.	Not Implemented (Superseded) - Image creation / recipes approach taken instead.
3.4	As a HN consumer, I can upload a new image to the cloud provider, such that I can instantiate this image following user story #1 and #2.	Implemented
3.5	As a HN consumer, I can save an existing VM as a reference image, such that I can instantiate this image following user story #1 and #2.	Implemented
3.6	As a HN consumer, I can transfer/convert an existing image from two cloud providers, using API calls, such that I do not have to recreate the image every time I change cloud provider.	Superseded

3.4 System Provisioning

Req Id	Requirement	Implementation State

4.1	As a HN consumer, I can provision a system (see definition above), in a single API call, such that the system is configured and ready to deliver the intended value. When required, the system provisioning can include parameters that the user must provide with the API call (e.g. external service endpoint, source of software, license information).	Implemented

3.5 Data Transfer and Placement

Req Id	Requirement	Implementation State
5.1	As a HN consumer, I can transfer a large amount of data (>TB?) in and out of a given cloud, such that I can push data into the cloud of my choice, or pull data out of a given cloud.	Not Implemented
	As a HN consumer, I can provision computing resources close to the data I am interested in processing, such that I do not have to transfer large amount of data to process it.	Not Implemented
	As a HN system administrator, I can easily replicate popular datasets in my cloud, such that users do not incur large data transfer costs to crunch the data on my cloud.	Not Implemented

3.6 Image Marketplace

Req Id	Requirement	Implementation State
6.1	As a HN consumer, I can access a catalogue of private and public images, such that I can easily search and discover them.	Implemented
6.1	As a HN consumer, I can refer to an image via a unique identifier, such that I can more easily know if this image is present in a given cloud or not. This means that images are annotated and/or tagged with unique identifier, and maintained across copy and replication processes.	Implemented
6.1	As a HN consumer, I can apply filters to my searches, such the following: · Architecture: e.g. 32 bit, 64 bit · O/S: e.g. Linux(all flavours), Windows, Solaris, AIX · Disk layout · Purpose, keywords · Installed software: e.g. package list · Cost · Capabilities and limitations in order to ease the selection process	Not Implemented
6.1	As a HN system administrator, I can access basic analytics, such	Not Implemented

	that I can see what images are popular and why.	
6.1	As a SME (for example), I can register and advertise in the marketplace images pre-installed and configured with added value software, such that I can generate revenue from these images. This implies a 'AWS devpay'-like feature, such that cloud providers can invoice users extra charges that get paid to the image owner.	Not Implemented

3.7 Security

Req Id	Requirement	Implementation State
7.1	As a HN consumer, I can manage firewall rules on vendor firewall systems protecting the cloud instantiated VM, such that I can control access of my cloud resources.	Not Implemented - Can be controlled at a Cloud Provider level
7.2	As a HN consumer, I can set up firewall rules between different Cloud Providers, in order to establish transparent networks.	Not Implemented
7.3	As a HN consumer, I can manage my instance keys, such that each have pre-assigned firewall rules, which I can specify when provisioning instances, or modify while the instances are running.	Not Implemented
7.4	As a HN consumer, I have a mechanism to rotate login credentials and instance keys, such that different instances are started with different credentials and keys, thus improving protection and security of my VMs.	Not Implemented
7.5	As a HN consumer, I can encrypt data in movement as well as persistent data.	Not Implemented
7.6	As a HN consumer, I can use the same login credentials to access all the HN services and Cloud Providers using SSO.	Not Implemented - Currently covered by a "key locker" approach meaning users only need to use one password to access the various clouds configured for their blue box account.

3.8 Networking

Req Id	Requirement	Implementation State
8.1	As a HN consumer, I can assume that all cloud providers are connected to GEANT, such that I have high-quality and predictable network performance between resources running in different HN cloud providers.	Partially Implemented

3.9 Cross Cloud / Federated Provisioning

Req Id	Requirement	Implementation State
9.1	As a HN consumer, I can provision cloud resources on several clouds at the same time, such that I can take advantage of the several clouds (e.g. geographical distribution, specific cloud provider feature and/or capabilities, data proximity).	Implemented

3.10 Hybrid Cloud Deployment

Req Id	Requirement	Implementation State
10.1	As a demand-side system administrator, I can install a Helix Nebula connector, such that my cloud provisioning solution can provision resources on the public HN cloud, alongside local cloud resources.	Implemented
10.2	As a demand-side system administrator, I can provision HN cloud resources able to connect to existing resources in my private cloud, such that a secure access is granted between the public HN resources and my private resources.	Implemented

3.11 Service Discovery

Req Id	Requirement	Implementation State
11.1	As a HN consumer, I can access the HN service catalog (e.g. service discovery service), such that I can discover what cloud provider is available, and what capabilities they each offer and support.	Partially Implemented
11.2	As a HN system administrator, I can populate/interconnect with the HN service catalog, such that I can advertise my cloud's capabilities, in order to attract users onto my cloud.	Partially Implemented

3.12 Data Marketplace

Req Id	Requirement	Implementation State
12.1	As a HN consumer, I can access a datamart (i.e. dataset marketplace), such that I can discover high-quality scientific	Not Implemented

	data, which I can process to generate high value assets.	
12.2	As a HN consumer, I can upload high-quality datasets to the datamart (i.e. dataset marketplace), such that I can share the datasets with the wider scientific community.	Not Implemented
12.3	As a HN consumer, I can protect the datasets I share on the HN federated cloud, such that only a restricted set of users (aka Virtual Organisations) have access to the data.	Not Implemented
12.4	As a HN consumer, I can refer to a dataset via a unique identifier, such that I can more easily know if this dataset is present in a given cloud or not. This means that datasets are annotated and/or tagged with unique identifier, and maintained across copy and replication processes.	Not Implemented
12.5	As a SME (for example), I can register and advertise in the marketplace added value datasets, such that I can generate revenue from these. This implies an 'AWS devpay'-like feature, such that cloud providers can invoice users extra charges that get paid to the dataset owner.	Not Implemented
12.6	As a HN system administrator, I can access basic analytics, such that I can see what datasets are popular and why.	Implemented

3.13 Analysis Summary

In general, the platform that is being rolled out for the Helix Nebula Marketplace (HNX) provides a good overall coverage of the requirements that have been identified in sub-groups Simple Provisioning, Contextualisation, Image Management and System Provisioning. However, at a Blue Box level, there are still further developments to be considered to cover Data Transfer and Placement and Security.

As it stands, little to no progress has been made the Data Marketplace activities beyond the efforts made specifically within the ESA flagship however there are efforts within the Helix Nebula consortium as a whole to identify ways in which the requirements of such a Data Marketplace can be implemented going forward.

4. User Requirements Analysis

The following section contains an analysis the requirements that were gathered as part of the requirements framework that was put in place in Work Package 3 of the project for provisioning, security, networking, billing and user level reporting. Together with the analysis presented in Section 3, a set of high-level future requirements formulated as user stories have been presented that should be considered for future implementation of the platform.

2.2 Technical Support Models

Within the Helix Nebula Marketplace production environment (HNX) a helpdesk tool, email support and phone support mechanism has been put in place, however as part of the user requirements analysis it was identified that users would also prefer mechanisms for Online Chat and a FAQ provision to be put in place.

Therefore the following requirements for the future have been identified

- REQ-TSM-1.1: As a user, I can seek assistance via an on-line chat with the helpdesk for Helix Nebula
- REQ-TSM-1.2 As a user, I can view a list of frequently asked questions via a web-site covering frequent technical and opportunity aspects of Helix Nebula

2.3 Provisioning

Provisioning of infrastructure via Helix Nebula is supported by both a web console and a fully documented API. However, an analysis of the findings of the flagship deployment activities found that the current blue box platform does not have a storage API in place that is required by some of the flagship deployments.

Therefore, the following three requirements from the initial Tech Arch requirements document will be carried forward:

- REQ-P-1.1: As a HN consumer, I can create, in a single API call, a new empty disk/volume, such that I can use it to build a new machine or as an additional disk/volume to a running VM (without the need of rebooting it)
- REQ-P-1.2: As a HN consumer, I can provision as for user story #1, with extra disks/volumes, such that I can keep my reference VM small and change available disk space using extra disks/volumes. The extra disks can be: persisted or volatile
- REQ-P-1.3: As a HN consumer, I can create, in a single API call, backup snapshots of my disks/volumes, such that I can re-provision and recover destroyed volumes

2.4 Security

The User Requirements Analysis from Work Package 3 identified that 100% of users would like to configure security via API and 75% would like to configure this via the web-console. Currently, Security cannot be configured directly via the Blue Box software, therefore the following requirements will be suggested for the roadmap of the brokerage platform:

- REQ-SEC-1.1: As a HN consumer, I can manage firewall rules on vendor firewall system on the cloud instantiated VM, such that I can control access of my cloud resources
- REQ-SEC-1.2: As a HN consumer, I can set up firewall rules between different Cloud Providers, in order to establish transparent networks
- REQ-SEC-1.3: As a HN consumer, I can manage my instance keys, such that each have pre-assigned firewall rules, which I can specify when provisioning instances, or modify while the instances are running
- REQ-SEC-1.4: As a HN consumer, I have a mechanism to rotate login credentials and instance keys, such that different instances are started with different credentials and keys, thus improving protection and security of my VMs
- REQ-SEC-1.5: As a HN consumer, I can encrypt data in movement as well as persistent data.

Additionally, the current Blue Box solution does not offer single-sign on via a single set of credentials. However, there is a "Key-Locker" approach in place that means users only need to make regular user of their Blue Box credentials as the credentials and keys for the various supported clouds are then configured and associated with their account in the platform. However, the solution should move towards a fully SSO solutions and therefore the following requirement is carried forward:

- REQ-SEC_1.6: As a HN consumer, I can use the same login credentials to access all the HN services and Cloud Providers using SSO

2.5 Networking

Currently the four supported cloud providers of Helix Nebula (Atos, CloudSigma, Interoute and T-Systems) are configured for GEANT connectivity, however the longer term support of GEANT for Helix Nebula must be clarified, therefore the following requirement is carried forwards:

- REQ-NET-1.1: As a HN consumer, I can assume that all cloud providers are connected to GEANT, such that I have high-quality and predictable network performance between resources running in different HN cloud providers

2.6 Billing

From the requirements analysis performed in Work Package 3 it is clear that users would like to access billing and metering information via both the API and Web Console. Currently the Blue Box does not support this and in addition the usage information is based upon consumed resources as opposed to the monetary value of those resources that are being consumed. The following requirements are proposed going forwards:

- REQ-BIL-1.1: As a HN consumer, I can view the costs incurred of the resources I am consuming within Helix Nebula across multiple providers
- REQ-BIL-1.2: As a HN consumer, I can access the costs incurred via the web console of the Blue Box
- REQ-BIL-1.3: As a HN consumer, I can access programmatically the resources currently being consumed via the API
- REQ-BIL-1.4: As a HN consumer, I can programmatically the cost of the resources that I am currently consuming via the Blue Box API.

2.7 Service Level Reporting

All users from which requirements were gathered identified requirements related to Service Level Reporting and

- REQ-SLR-1.1: As a HN consumer, I can access service level reports related to the service levels achieved during a specified period across one or multiple providers as configured by the user
- REQ-SLR-1.2: As a HN consumer, I can programmatically access via an API the service level reports
- REQ-SLR-1.3: As a HN consumer, I can access via a web-console service level reports

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

5. Other Requirements

The following section contains a set of User Stories/Requirements that were not explicitly covered by the analysis performed in WP3, but were contained in the User Stories analysis presented as part of the Tech Arch document. These are requirements that have not yet been addressed by the platform in place today.

2.8 Data Management

The following Data Management requirements should be carried forward:

- REQ-DM-1.1: As a HN consumer, I can transfer a large amount of data (>TB?) in and out of a given cloud, such that I can push data into the cloud of my choice, or pull data out of a given cloud.
- REQ-DM-1.2: As a HN consumer, I can provision computing resources close to the data I am interested in processing, such that I do not have to transfer large amount of data to process it.
- REQ-DM-1.3: As a HN system administrator, I can easily replicate popular datasets in my cloud, such that users do not incur large data transfer costs to crunch the data on my cloud.

2.9 Image Marketplace

The following Image Management requirements have not yet been addressed by the solution in place and should be carried forward:

REQ-IM-1.1: As a HN consumer, in order to ease the selection process, I can apply filters to my searches, such the following:

- Architecture: e.g. 32 bit, 64 bit
- O/S: e.g. Linux(all flavours), Windows, Solaris, AIX
- Disk layout
- Purpose, keywords
- Installed software: e.g. package list
- Cost
- Capabilities and limitations

- REQ-IM-1.2: As a HN system administrator, I can access basic analytics, such that I can see what images are popular and why
- REQ-IM-1.3: As a SME (for example), I can register and advertise in the marketplace images pre-installed and configured with added value software, such that I can generate revenue from these images. This implies an 'AWS devpay'-like feature, such that cloud providers can invoice users extra charges that get paid to the image owner.

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

2.10 Data Marketplace

The following requirements were not considered in-scope for the first release of the Helix Nebula Marketplace, however, have always been considered part of the longer term goals of Helix Nebula:

- REQ-DMP-1.1: As a HN consumer, I can access a datamart (i.e. dataset marketplace), such that I can discover high-quality scientific data, which I can process to generate high value assets
- REQ-DMP-1.2: As a HN consumer, I can upload high-quality datasets to the datamart (i.e. dataset marketplace), such that I can share the datasets with the wider scientific community
- REQ-DMP-1.3: As a HN consumer, I can protect the datasets I share on the HN federated cloud, such that only a restricted set of users (aka Virtual Organisations) have access to the data
- REQ-DMP-1.4: As a HN consumer, I can refer to a dataset via a unique identifier, such that I can more easily know if this dataset is present in a given cloud or not. This means that datasets are annotated and/or tagged with unique identifier, and maintained across copy and replication processes.
- REQ-DMP-1.5: As a SME (for example), I can register and advertise in the marketplace added value datasets, such that I can generate revenue from these. This implies a 'AWS devpay'-like feature, such that cloud providers can invoice users extra charges that get paid to the dataset owner.
- REQ-DMP-1.6: As a HN system administrator, I can access basic analytics, such that I can see what datasets are popular and why.

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301

6. Conclusion

In summary, the analysis performed in this document identifies at a high level the following items that need to be addressed going forward in order to further align Helix Nebula with the requirements of the user community currently making use of the platform:

- Provisioning: Users can successfully provision virtual resources via both the web-console and API, however further iteration is required in the form of support for a Storage and Disk Management
- Security: Currently there is no support for security configuration via the Blue Box and a series of requirements have been identified going forward
- Image Marketplace: A basic service catalogue has been introduced within the Blue Box but requirements have been identified to support better searching and selection of the services available within Helix Nebula
- Data Management: Data Management requirements have to date not been addressed but are contained on the roadmap for Helix Nebula
- Data Marketplace: The Data Marketplace requirements have not yet been addressed however a H2020 proposal has been submitted called XzelCloud that will address these particular needs
- Billing & Service Level Reporting: Evolution of the Billing and Metering functionality will be required in order to address the needs of the users and in addition programmatic access to Service Level Reporting is required.

The Helix Nebula project is co-funded by the European Community Seventh Framework Programme (FP7/2007-2013) under Grant Agreement no 312301