Shapes in the Cloud: Defining Unidata's Efforts

Strategic Advisory Committee

October 2014

Ward Fisher

Goal

- Provide context for the Cloud-related work Unidata is doing.
- Discuss the status of our current ongoing Projects.

Context

- “The Cloud” can have different meanings to different organizations.
- When Unidata talks about ‘The Cloud’, what do we mean?

Using the Cloud

The Cloud for Data Storage

Using the Cloud

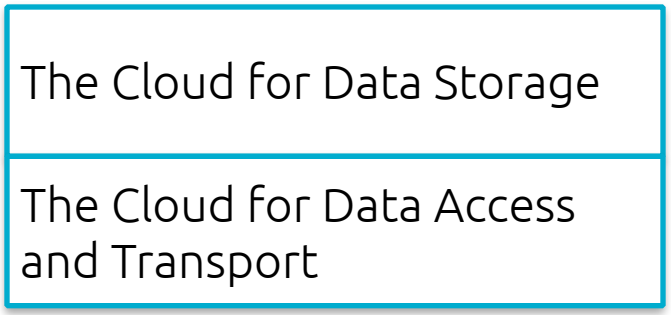

Using the Cloud

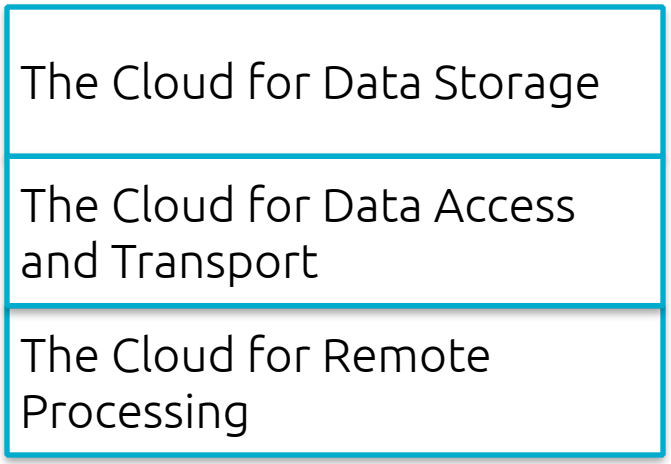

Uses for the Cloud

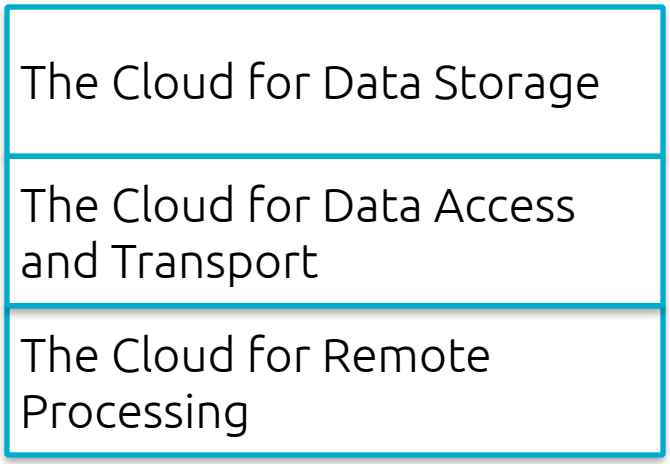

Uses for the Cloud

The Cloud for Data Storage

The Cloud for Data Access and Transport

The Cloud for Remote Processing

Harnessing the Cloud

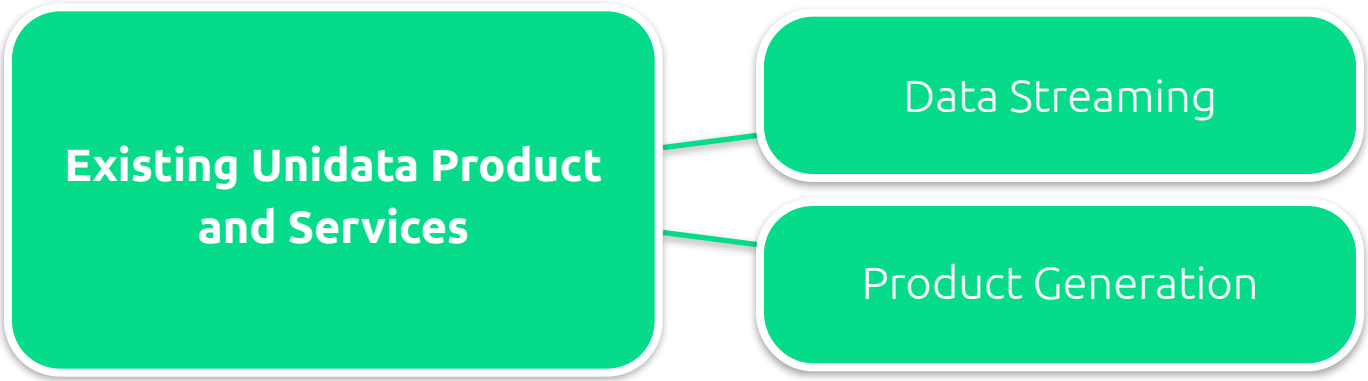

Harnessing the Cloud

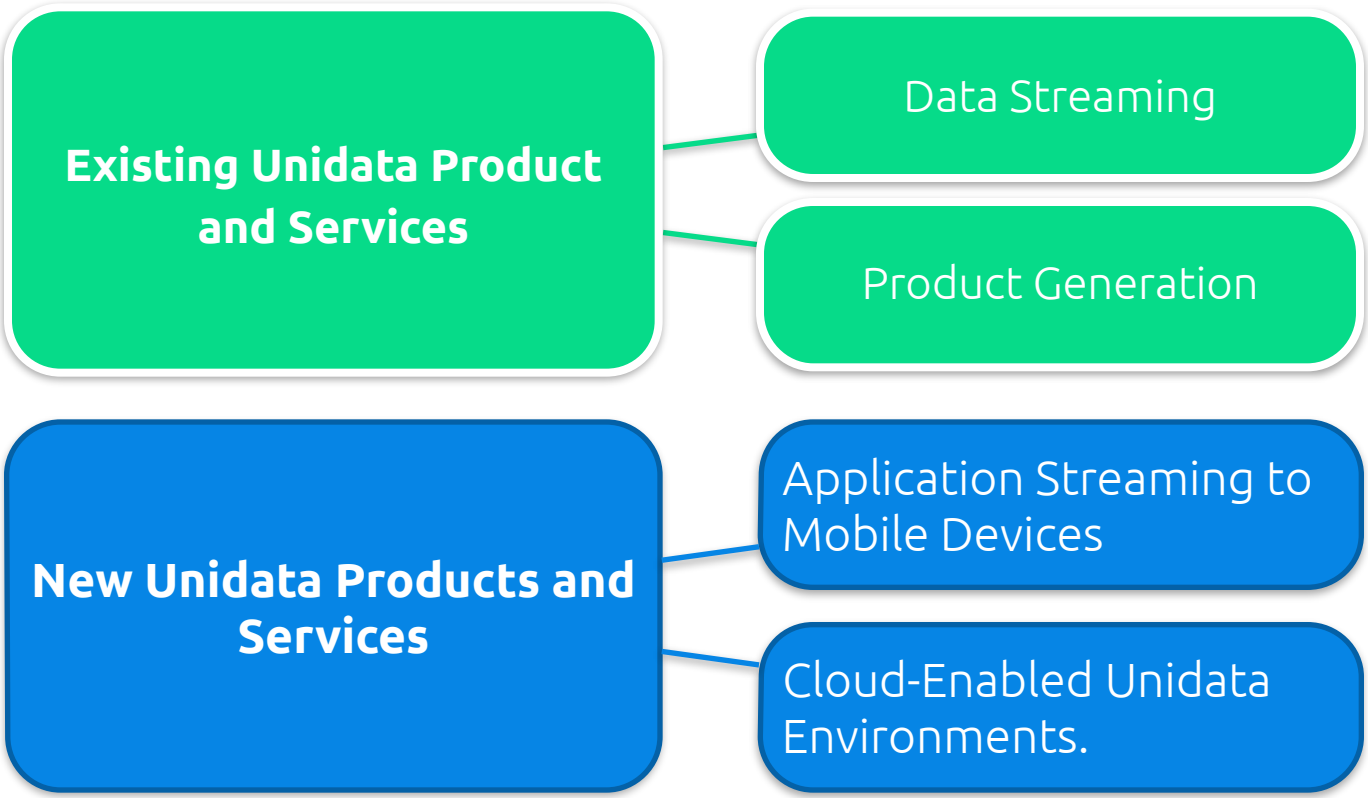

Ongoing Projects

- AWIPS II Cloud Servers

Ongoing Projects

- AWIPS II Cloud Servers
- IDD Product Generation and Additional Experimentation

Ongoing Projects

- AWIPS II Cloud Servers
- IDD Product Generation and Additional Experimentation
- “Unidata-in-a-Box” Virtual Environment.

Ongoing Projects

- AWIPS II Cloud Servers
- IDD Product Generation and Additional Experimentation
- “Unidata-in-a-Box” Virtual Environment.
- IDV Application-Streaming Cloud Servers

AWIPS II Cloud Servers

- Unidata is testing small footprint **EDEX** servers (no **NEXRAD** Level 2 or 3 or high-resolution **CONDUIT** models) on both **Microsoft Azure** and **Amazon EC2** cloud server environments.

AWIPS II Cloud Servers

- EC2 Instance is created cooperatively with Embry Riddle Aeronautical University (ERAU) as part of their equipment grant award.

AWIPS II Cloud Servers

- The Azure instance is serving data to AWIPS II 14.2.1 beta testers.

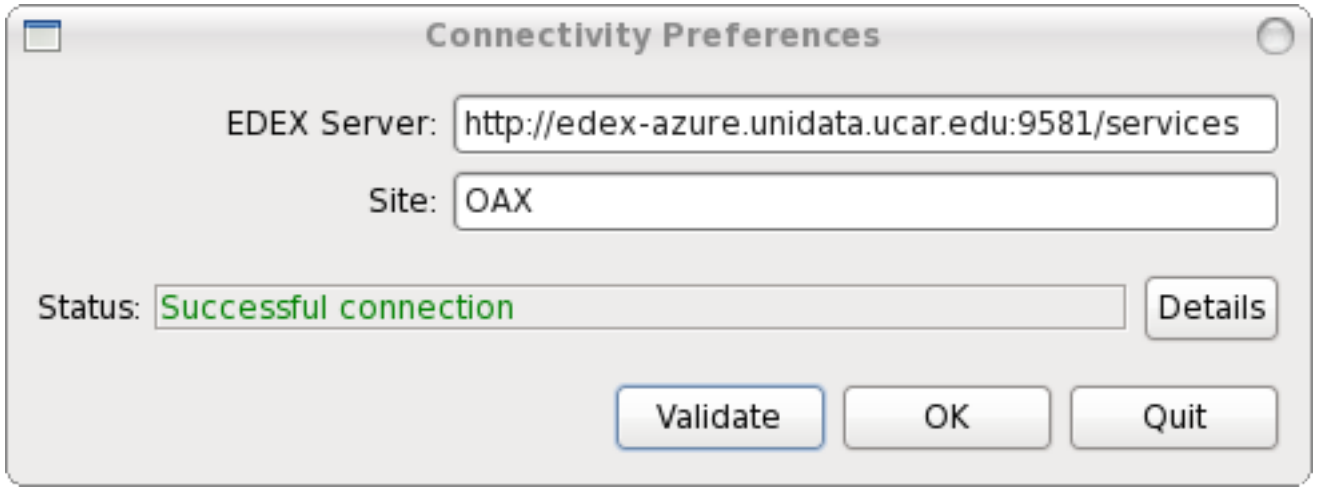

IDD Product Generation and Additional Experimentation

- Unidata operates mid-sized instances in **Azure** and **Amazon EC2** clouds.
- These instances are being used to generate image products for the **IDD FNEXRAD** and **UNIWISC** data streams.

IDD Product Generation and Additional Experimentation

- EC2 Instance is the primary source of FNEXRAD and UNIWISC data streams to IDD participants.
- We will be transitioning to Azure cloud instances to reduce recurring costs, due to an resource award from Microsofts **Azure-for-Research** program.

IDD Product Generation and Additional Experimentation

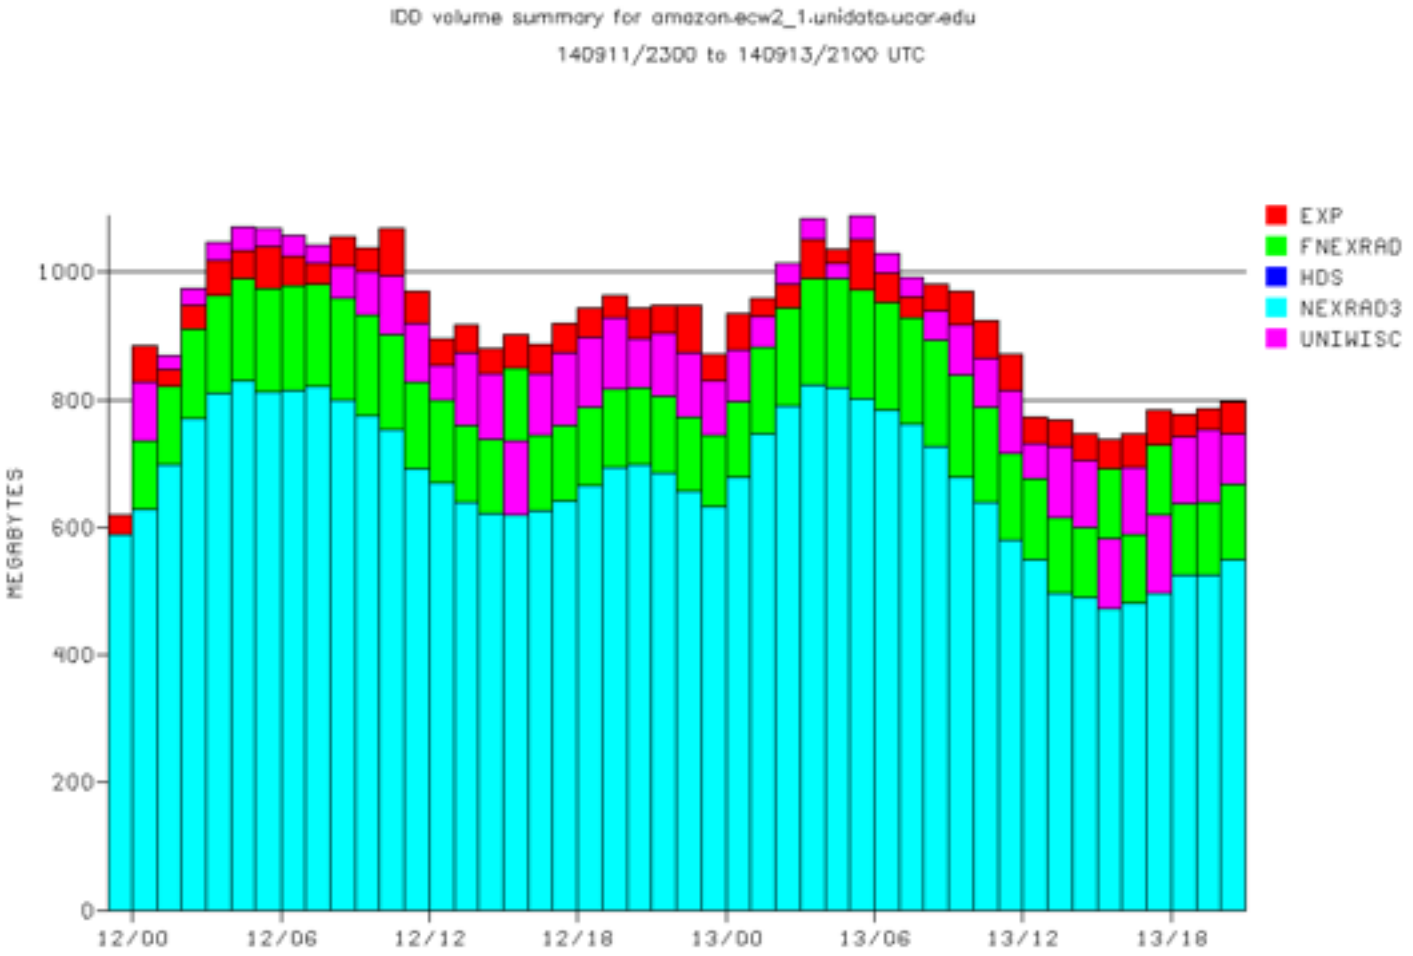

IDD Product Generation and Additional Experimentation

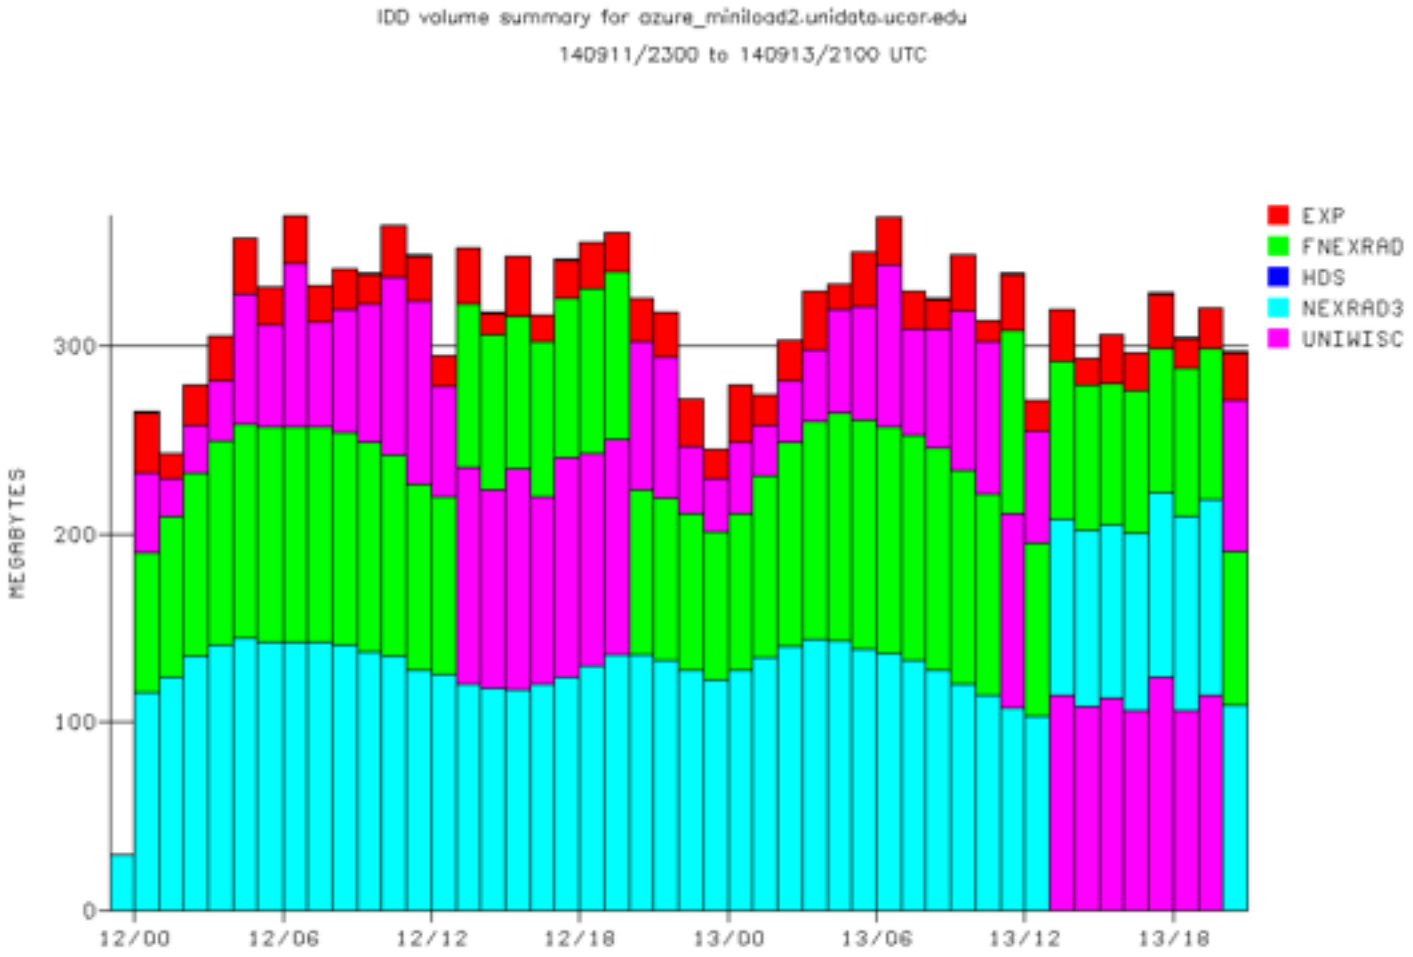

New Cloud-Enabled Projects

- Unidata-in-a-Box Virtual Environment.
- Integrated Data Viewer (IDV) on mobile devices via Application Streaming.

Unidata-in-a-Box

- This is an effort to provide a VM image pre-configured with Unidata software and services.

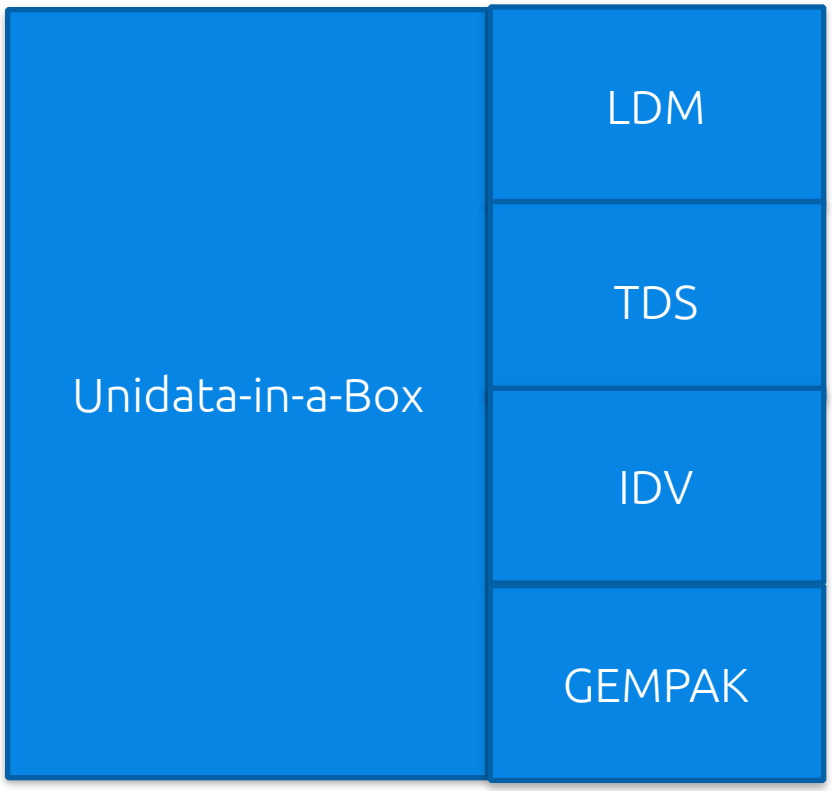

Unidata-in-a-Box

- This will be delivered via cloud services like Vagrant, or through traditional file transfer.

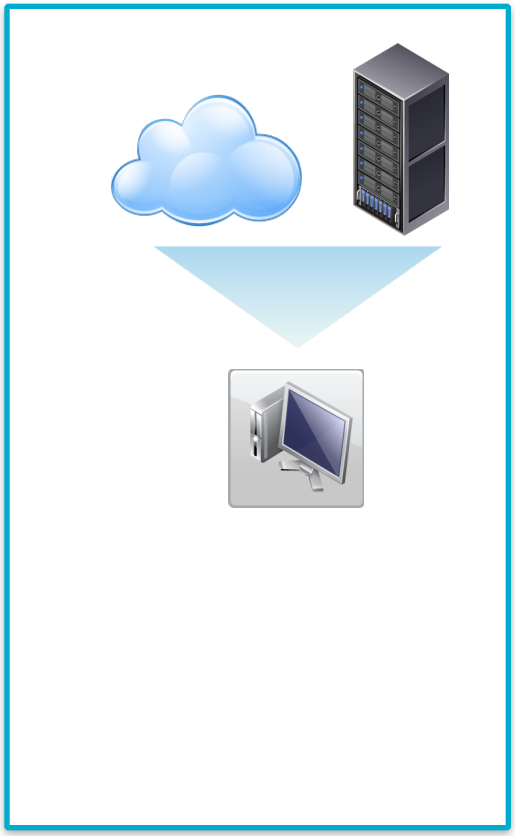

Unidata-in-a-Box

- UiaB can then be deployed within a sandboxed environment on the local PC.

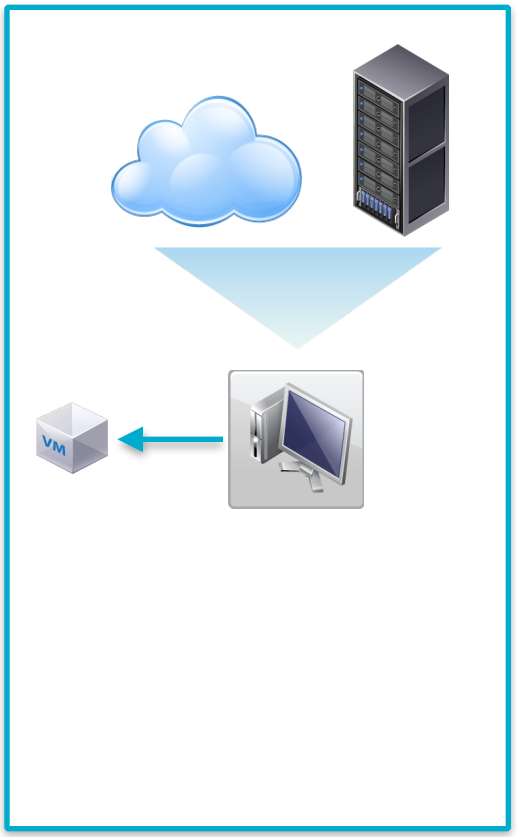

Unidata-in-a-Box

- It can also be deployed back into the Cloud.

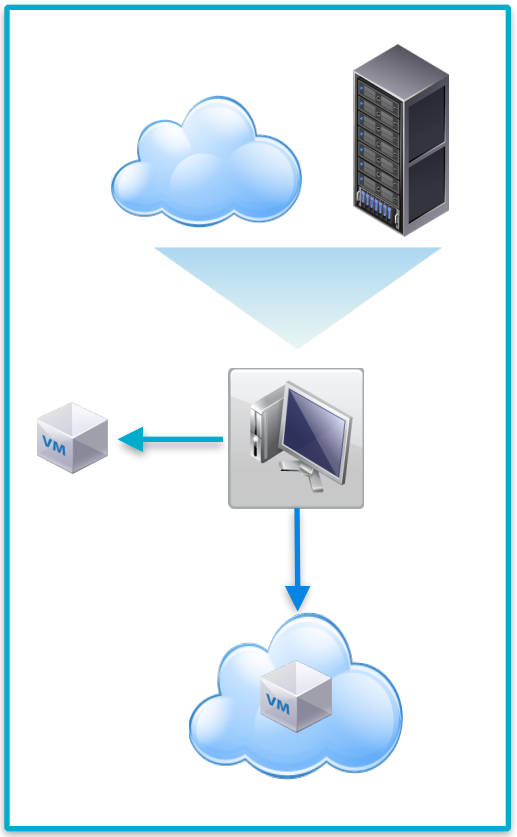

Unidata-in-a-Box

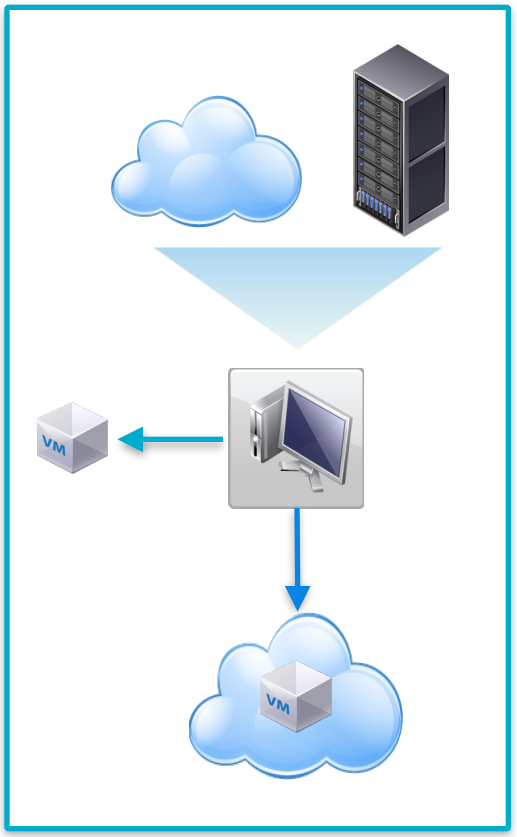

Unidata-in-a-Box

- Status: Ongoing

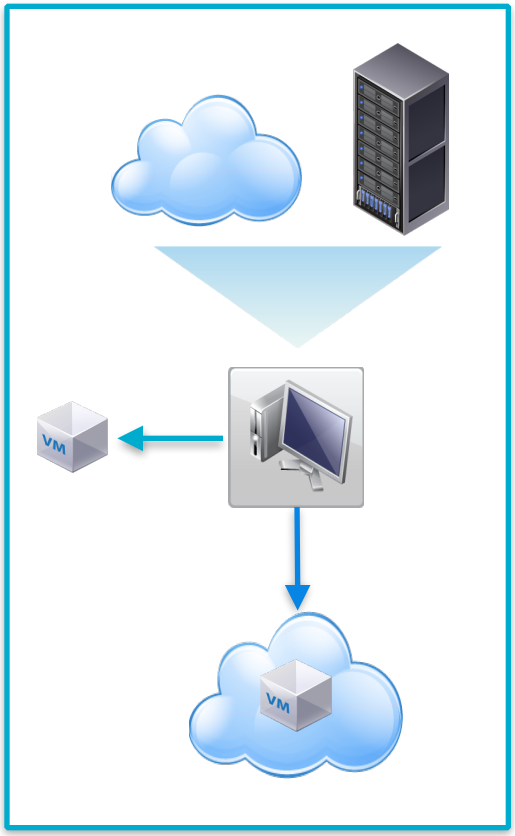

IDV via Application Streaming

- Goal: Create and provide IDV instances which live in the cloud but may be *streamed* to various devices.

IDV via Application Streaming

- Goal: Create and provide IDV instances which live in the cloud but may be *streamed* to various devices.
- Drawback: The interface is adapted to the target device but is not optimized for it.

IDV via Application Streaming

- Goal: Create and provide IDV instances which live in the cloud but may be *streamed* to various devices.
- Drawback: The interface is adapted to the target device but is not optimized for it.
- Benefit: Brings the IDV to new classes of devices without needing to modify the IDV.

Application Streaming?

- Application Streaming is similar to **remote desktop** technology, but is meant to stream a single application.

Application Streaming?

- Application Streaming is similar to **remote desktop** technology, but is meant to stream a single application.
- The server instance is optimized for the dimensions of the remote client device.

Status

- Using the Azure Web API, we are able to dynamically allocate and provision VMs used to host individual IDV instances.

Status

- Using the Azure Web API, we are able to dynamically allocate and provision VMs used to host individual IDV instances.
- We are then able to instantiate IDV instances then streamed (via existing remote-desktop protocols) to mobile devices.

Next Step

- Current efforts are focused on creating a web dashboard which will allow users to register and manage IDV-streaming requests.

Performance

- How well does it perform?

Performance

- How well does it perform?
- Performance is tied to the client used.

Performance

- How well does it perform?
- Performance is tied to the client used.
  - Dedicated clients such as "Parallels Remote Access" or "Air Login": Very good, typically adapted to touch interfaces.

Performance

- How well does it perform?
- Performance is tied to the client used.
  - Dedicated clients such as "Parallels Remote Access" or "Air Login": Very good, typically adapted to touch interfaces.
  - Generic VNC clients: acceptable, but suffer from inconsistent interfaces between clients.

Parallels Access Demo

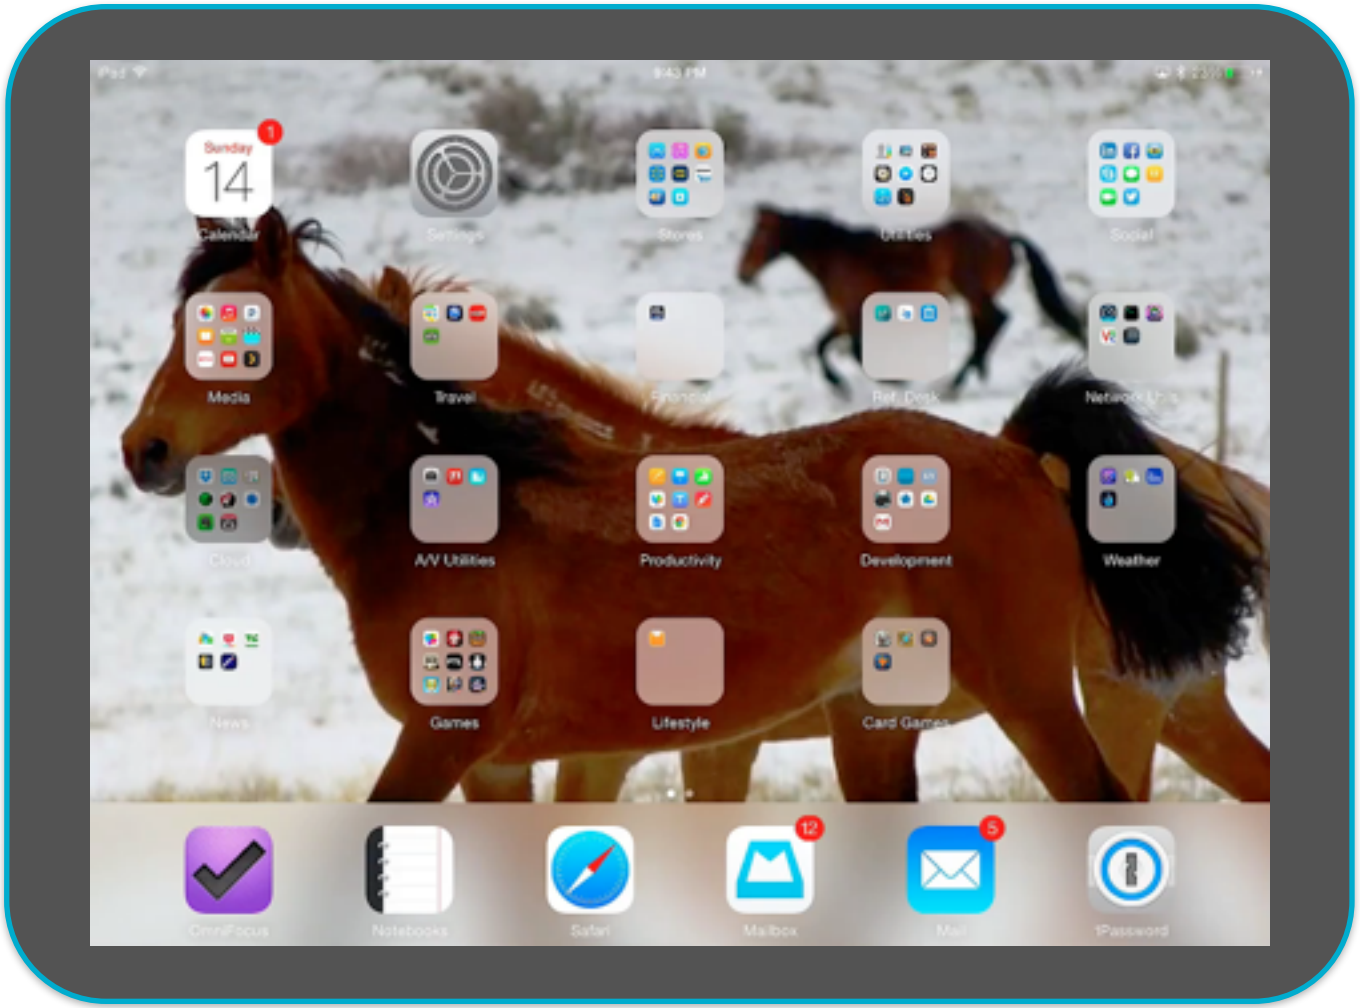

Parallels Access Demo

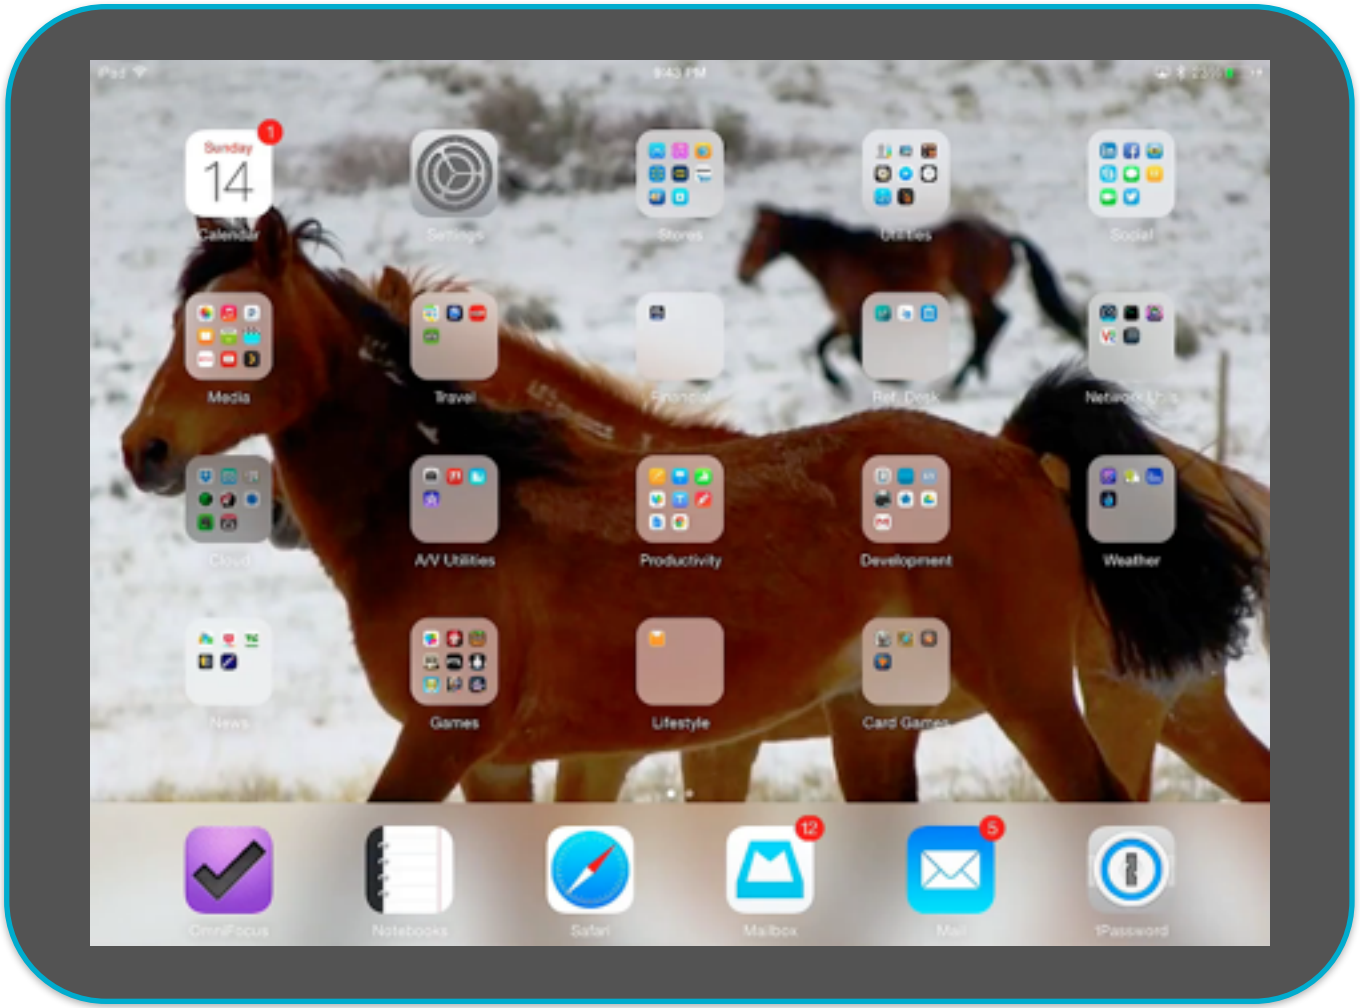

Future Work

- As a first attempt, the results have been very promising.

Future Work

- As a first attempt, the results have been very promising.
- Moving forward: generic VNC access or dedicated client access?

Future Work

- As a first attempt, the results have been very promising.
- Moving forward: generic VNC access or dedicated client access?
- The latter would be preferable given infinite resources, but we have not been given infinite resources (yet).

Internal Tools

- Github: Revision control, issue tracking, collaboration.

Internal Tools

- Github: Revision control, issue tracking, collaboration.
- Binstar: Binary python packages.

Internal Tools

- Github: Revision control, issue tracking, collaboration.
- Binstar: Binary python packages.
- CDash: Dashboard for software testing.

Internal Tools

- Vagrant: Cloud-enabled VM management for developers & scientists.

Internal Tools

- Vagrant: Cloud-enabled VM management for developers & scientists.
  - NetCDF testing exists inside VMs deployed by vagrant.

Internal Tools

- Vagrant: Cloud-enabled VM management for developers & scientists.
  - NetCDF testing exists inside VMs deployed by vagrant.
    - Git repository: http://github.com/WardF/tiny-ci

Summary

- Unidata is proceeding into the cloud along multiple fronts.

Summary

- Unidata is proceeding into the cloud along multiple fronts.
- These projects are no longer speculative.

Summary

- Unidata is proceeding into the cloud along multiple fronts.
- These projects are no longer speculative.
- What we learn from these projects will inform the direction of future projects.

Summary

- Unidata is proceeding into the cloud along multiple fronts.
- These projects are no longer speculative.
- What we learn from these projects will inform the direction of future projects.
- Recording an iPad screencast is difficult.

Questions?