

Roswell OnDemand

A One-Stop Shop for High Performance Computing



L. Shawn Matott, PhD
Senior Systems Analyst
IT Research Computing

Computing Assets

- Roswell HPCC
 - 1,600 processors, 450 TB Lustre storage
 - Software modules include: Grace, EGSnrc, Keras, and TensorFlow, many bio-informatics packages
 - OnDemand web interface
<http://u1.roswellpark.org/pun/sys/dashboard/>
 - e-mail ServiceDesk@RoswellPark.org to request access



A look inside Roswell's HPC data center.

A typical HPC Engagement

Discovery
Initial meeting between
research group and IT RC.

Preparation
IT RC performs S/W installation
and template development to
meet needs of research group.

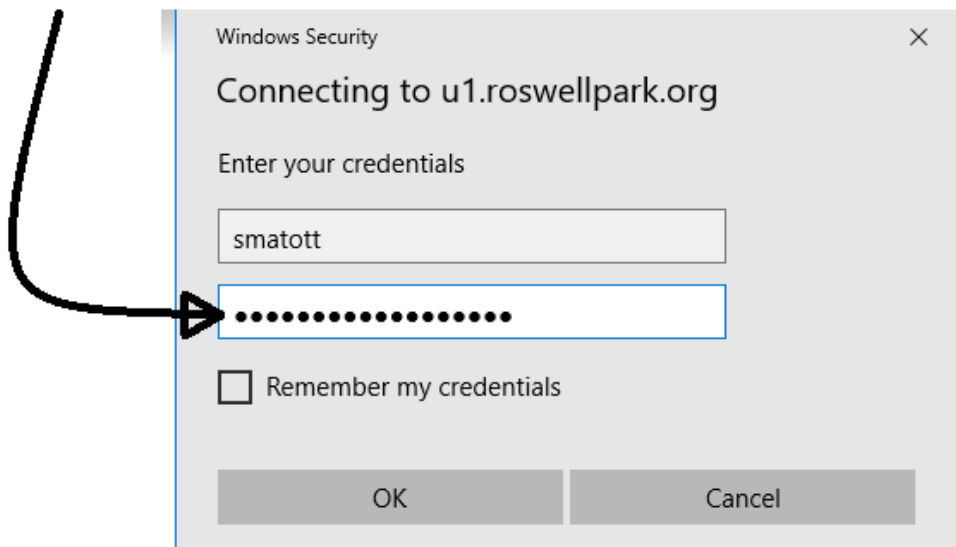
Setup/Training
Account creation and
training of new users.

Data Migration
Research group migrates
data into HPC storage.

Ongoing Research
Ongoing usage of
compute and storage
assets with periodic IT
support requests.

OnDemand Login

<http://u1.roswellpark.org/pun/sys/dashboard/>



- Note: *your OnDemand user ID and password may not be the same as your main Roswell ID and password*

OnDemand Dashboard


- Files
 - Upload, download, edit, move, delete etc.
- Jobs
 - Configure, submit, & monitor jobs
- Clusters
 - Linux CLI
- Interactive Apps
 - Launch programs with GUIs



The screenshot shows a web browser window with the title 'Dashboard' and a single tab. The address bar displays the URL <https://u1.roswellpark.org/pun/sys/dashboard/>. The dashboard header is a dark blue bar with the text 'Roswell Park HPC OnDemand' and navigation links for 'Files', 'Jobs', 'Clusters', 'Interactive Apps', and 'My Interactive Sessions'. Below the header is a yellow notification box with the text: 'NOTICE: The OnDemand interface is currently under development as of 2/8/19.' The main content area features the Roswell Park Comprehensive Cancer Center logo, which consists of a blue square with white radiating lines and the text 'ROSWELL PARK' in large black letters, with 'COMPREHENSIVE CANCER CENTER' in smaller blue letters below it. Underneath the logo, a line of text states: 'OnDemand provides an integrated, single access point for all of your HPC resources.' At the bottom of the dashboard, the text 'Message of the Day' is visible.

OnDemand - Files

- File Explorer
 - Transfers via upload/download or drag and drop
- Editor
 - Language aware coloring (e.g. Python comments, etc.)



OnDemand – Compute Jobs

- Composer
 - Create and configure a new job
 - Use templates to minimize effort
 - Monitor or alter current jobs
- Active Jobs
 - Check job queues

The screenshot displays the 'Job Composer' interface for Roswell Park HPC OnDemand. The browser address bar shows the URL <https://u1.roswellpark.org/pun/sys/myjobs>. The navigation bar includes 'Roswell Park HPC OnDemand / Job Composer' and tabs for 'Jobs' and 'Templates'. The main heading is 'Jobs'. Below this, there are buttons for '+ New Job', 'Edit Files', 'Job Options', 'Open Terminal', 'Submit', 'Stop', and 'Delete'. A search bar is also present. The job list table shows three completed jobs:

Created	Name	ID	Cluster	Status
April 25, 2019 10:47am	BiocParallel_multicore	736401	Roswell HPC Cluster	Completed
March 20, 2019 1:36pm	hello	728989	Roswell HPC Cluster	Completed
February 22, 2019 11:20am	Keras_TensorFlow_CIFAR10	725194	Roswell HPC Cluster	Completed

OnDemand – Job Templates

- Templates
 - Searchable
 - IT RC (i.e. Shawn) can create new templates on request
 - Users can also create and manage their own templates

The screenshot shows a web browser window with two tabs: 'Dashboard' and 'Job Composer'. The address bar shows the URL <https://u1.roswellpark.org/pun/sys/myjobs/workflows/new>. The page title is 'Templates'. Below the title, a message states: 'To create a new job, select a template to copy, fill out the form to the right, and click "Create New Job".'.

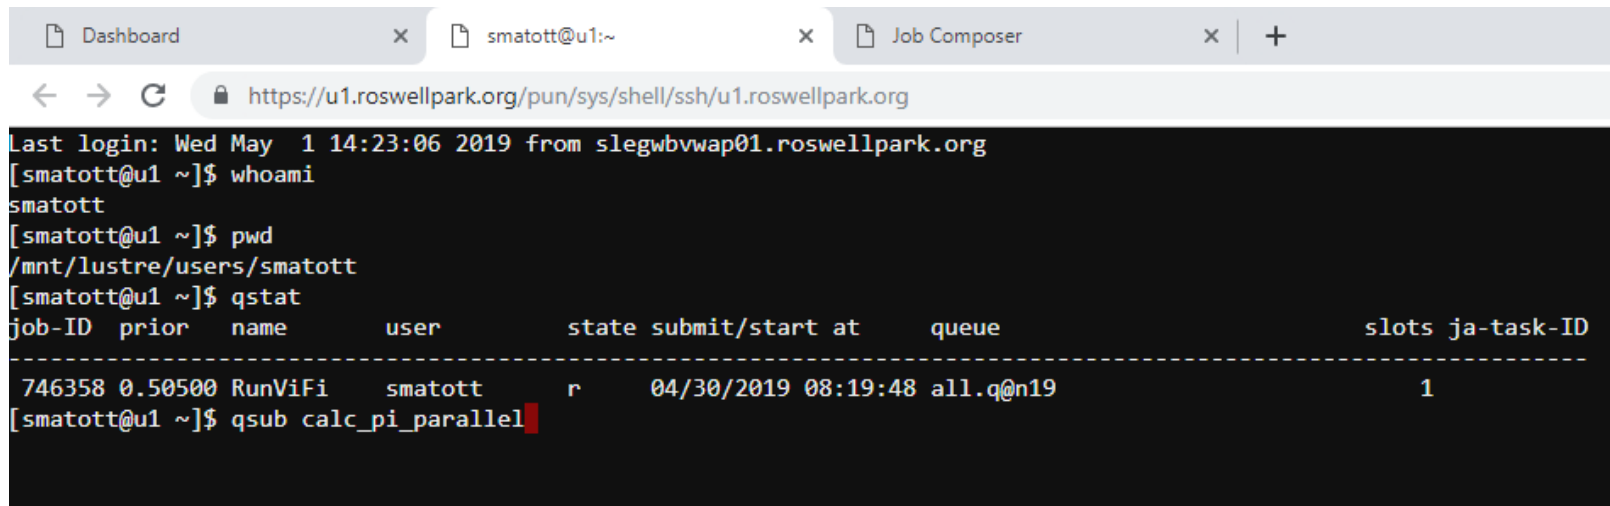
At the top of the main content area, there are two buttons: '+ New Template' and 'Copy Template'. Below these, there is a section for managing templates. It includes a 'Show' dropdown set to '10' and a 'Search' input field. A table lists the available templates:

Name	Cluster	Source
BlocParallel_multicore	Hpc	System Templates
BlocParallel_Rmpi	Hpc	System Templates
BlocParallel_snow	Hpc	System Templates
hello	Hpc	System Templates
hello_array	Hpc	System Templates
Keras_TensorFlow_charseq	Hpc	System Templates

Below the table, there is a form titled 'Create New "BlocParallel_multicore"'. The form includes a description: 'This illustrates using BlocParallel to parallelize a set of "busy" calculations across processors on a single node.' and three fields: 'Job Name' (containing 'BlocParallel_multicore'), 'Cluster' (a dropdown menu showing 'Roswell HPC Cluster'), and 'Script Name' (containing 'sge-bloc-parallel-multicore'). At the bottom of the form are two buttons: 'Create New Job' and 'Reset'.

OnDemand – Cluster CLI

- Provides a bash command line interface on front-end node (u1)
- Handy for those who are familiar with conventional HPC access paradigm (e.g. putty, ssh)



```
Last login: Wed May 1 14:23:06 2019 from slegwbvwap01.roswellpark.org
[smatott@u1 ~]$ whoami
smatott
[smatott@u1 ~]$ pwd
/mnt/lustre/users/smatott
[smatott@u1 ~]$ qstat
job-ID prior  name      user      state submit/start at    queue      slots ja-task-ID
-----
746358 0.50500 RunViFi    smatott   r       04/30/2019 08:19:48 all.q@n19      1
[smatott@u1 ~]$ qsub calc_pi_parallel
```

OnDemand – Interactive Apps

- Special compute jobs for launching interactive GUI-based applications
- Remote Desktop + Individual Apps
- OpenGL codes can benefit from GPU hardware acceleration (viz nodes)

The screenshot shows the Roswell Park HPC OnDemand web interface. The browser address bar displays the URL: https://u1.roswellpark.org/pun/sys/dashboard/batch_connect/sys/igv/session_contexts/new. The navigation bar includes links for Files, Jobs, Clusters, Interactive Apps, and My Interactive Sessions. A yellow notice banner states: "NOTICE: The OnDemand interface is currently under development as of 2/8/19." Below the navigation bar, a breadcrumb trail shows Home / My Interactive Sessions / IGV. The main content area features a sidebar with a list of Interactive Apps: Desktops, Remote Desktop, Programs, Grace, IGV (selected), Jupyter, ParaView, RGui, and RStudio. Below this is a section for Interactive Apps [Sandbox] with a Programs list containing xclock. The main panel displays the IGV app configuration. It includes a description: "This app will launch IGV - the Integrative Genomics Viewer." and a section for IGV Version, where 2.5.0 is selected from a drop-down menu. Below this is an Account field, a Queue field set to all.q, a Number of hours field set to 1, and a Number of cores field set to 1. At the bottom, there is a checkbox for "I would like to receive an email when the session starts" and a prominent blue Launch button.

← → ↻ https://u1.roswellpark.org/pun/sys/dashboard/batch_connect/sys/igv/session_contexts/new

Roswell Park HPC OnDemand Files Jobs Clusters Interactive Apps My Interactive Sessions

NOTICE: The OnDemand interface is currently under development as of 2/8/19.

Home / My Interactive Sessions / IGV

Interactive Apps

Desktops

Remote Desktop

Programs

Grace

IGV

Jupyter

ParaView

RGui

RStudio

Interactive Apps [Sandbox]

Programs

xclock

IGV

This app will launch IGV - the Integrative Genomics Viewer.

IGV Version

2.5.0

Please select a version of IGV from the drop-down.

Account

Queue

all.q

Number of hours

1

Number of cores

1

☐ I would like to receive an email when the session starts

Launch

Use the app window at left to select an app and submit request via “Launch” button.

OnDemand – Interactive Sessions

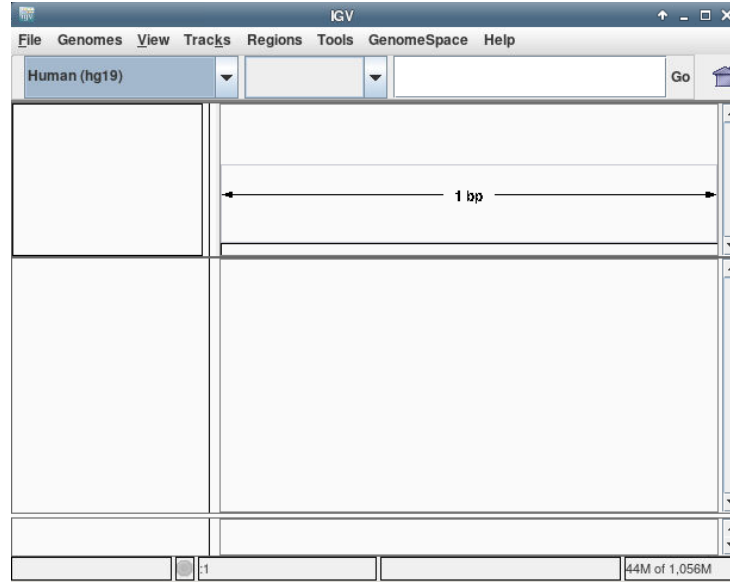
- Interactive apps are compute jobs!
- Can close browser window and reconnect without loss of work.

The screenshot displays the Roswell Park HPC OnDemand web interface. The browser address bar shows the URL: https://u1.roswellpark.org/pun/sys/dashboard/batch_connect/sessions. The navigation bar includes links for Files, Jobs, Clusters, Interactive Apps, and My Interactive Sessions. A yellow notice banner states: "NOTICE: The OnDemand interface is currently under development as of 2/8/19." Below this, a green message box confirms: "Session was successfully created." The main content area shows a sidebar with a list of interactive apps: Remote Desktop, Grace, IGV, and Jupyter. The IGV app is selected, showing details for session 746925, which is running on 1 core. The session was created on 2019-05-02 at 15:37:15 EDT. A "Launch IGV" button is visible, along with a "Delete" button and a "View Only (Share-able Link)" button.

Once the job starts the window will refresh and offer a “Launch [app]” button. Click it to open the app in a browser window.

OnDemand – Interactive Sessions

- Interactive apps are compute jobs!
- Can close browser window and reconnect without loss of work.



In this example, a remote desktop is opened and pre-populated with an instance of the IGV genome viewer.

OnDemand – Machine Learning example

- Dashboard → Jobs → Job Composer → Templates
- Search for “Keras”
- Select the “wordseq” template
 - Generate HPL-style prose!
- Click “Create New Job”
 - Drops you into the “Jobs” interface
- Explore the job details area
- Click on the “.py” file to open in editor
 - contains the Keras code
- Find the section marked “define model”
 - Feel free to add, delete, modify layers

```
170     # define model
171     model = Sequential()
172     model.add(Embedding(vocab_size, 50, input_length=seq_length))
173     model.add(LSTM(100, return_sequences=True))
174     model.add(LSTM(100))
175     model.add(Dense(100, activation='relu'))
176     model.add(Dense(vocab_size, activation='softmax'))
177     print(model.summary())
178
179     # train the model
180     model.compile(loss='categorical_crossentropy', optimizer='ada
181     model.fit(X, Y, batch_size=128, epochs=100, verbose=2)
182
183     # calculate and report performance of the model
184     print("")
185     print("-----")
186     print_confusion_matrix(model, X, Y, inv_mapping)
187     print("-----")
188
189     # save model
190     model.save('model.h5')
```

OnDemand – Machine Learning example

- Find the section marked “procedure == “generate”
 - Feel free to adjust the generator sequence and see what sort of Lovecraftian text you can get the AI to come up with!
 - The sequence must be at least 10 words and include a space at the end.

```
322 ▾ elif(procedure == "generate"):
323     passage = generate_sequence("the horror of it all was quite something to behold ", 500)
324     passage = " ".join(str(x) for x in passage)
325     passage = passage.replace('!', '\n')
326     print("")
327     print("***20)
328     print(passage)
329     print("***20)
330     print("")
```

OnDemand – Machine Learning example

- Open the sge-keras-wordseq script in the editor
- Add lines at the bottom to run test and generate portions of the model:

```
23 python keras_wordseq_dunwich.py $NSLOTS train
24 python keras_wordseq_dunwich.py $NSLOTS test
25 python keras_wordseq_dunwich.py $NSLOTS generate
```

- Save changes
- Job Composer → Submit
- Wait for job to complete
- Open wordseq.out →

```
4130 *****
4131 the horror of it all was quite something to behold the panic bearing in household careful the copy
4132 muscular hellish elder and boarded non users
4133 and menaces and demoniac
4134 having was instruct of settled business in the fetor in all and grinned were this time wilbur afterward
4135 bits of no direction like night and only a restraint and very stigmata of loading of the neighboring
4136 and drag it and foul and a chance of putting ye the men then those fragments and and a hastily
4137 had less road there was
4138 crumpled to the
4139
4140 edition died came his use it seemed to comment in the wildest green allowed
4141 blasted albino yog sothoth he was in any use
4142 rate and way to certain available the singularly state of crushed lantern sounds as haow whose heads
4143 grow ahead
4144 seemed to let to seth is no frye then against any books fer
4145 aout over what was armitage walk from the hill indoors as the doors of growing void which even earl
4146 though pull he feared
4147 disappeared near down
4148 dogs he shuddered were tarlike ancient to let in wilbur was tell it was creators other distaste no
4149 only the bending beside the smell are
4150 why partly
4151 but in august he caved
```