

# Job Step Viewer

A new utility for working with  
`jsrun` on Summit

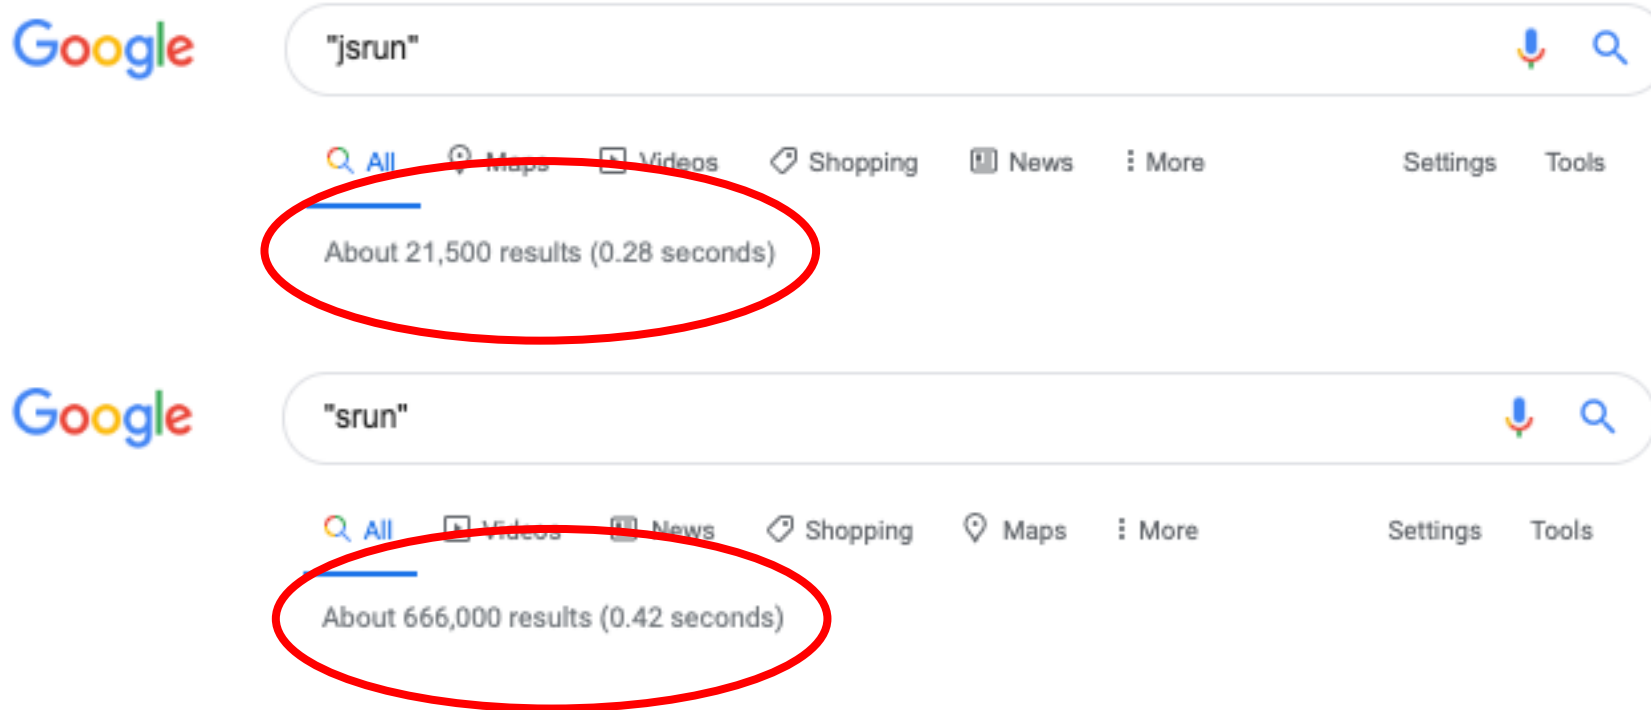
Jack Morrison  
HPC Engineer, User Assistance  
March 2020

ORNL is managed by UT-Battelle, LLC for the US Department of Energy

# What is Job Step Viewer?

- A new tool to generate graphical views of an application's runtime layout on Summit
  - Resource Sets
  - Process placement
  - Thread binding
  - Node locality
  - Sanity checking
- A replacement for jsrunVisualizer (<https://jsrunvisualizer.olcf.ornl.gov/>)
- <https://jobstepviewer.olcf.ornl.gov/>

# jsrun is (still) new



- Generates a lot of questions

# Goals for Job Step Viewer

1. Help end-users understand how `jsrun` options interact with their environment to launch applications on Summit
2. Lower the learning curve for new Summit users
3. Provide project teams and OLCF staff a way to discuss job steps
4. Encourage full utilization of Summit's large multi-GPU nodes



# Features

- Easy to use
- Stand-alone jsrun testing
- Informed by real runs on Summit & respects user's runtime environment
- Warnings for common mistakes
- SMT{1,2,4} support
- NVMe (Burst Buffer) support
- Supports Explicit Resource Files (ERF)
- Easily sharable, unique URLs

# Features

- Easy to use
- Stand-alone `jsrun` testing
- Informed by real runs on Summit & respects user's runtime environment
- Warnings for common mistakes
- SMT{1,2,4} support
- NVMe (Burst Buffer) support
- Supports Explicit Resource Files (ERF)
- Easily sharable, unique URLs

```
login1.summit$ bsub -W 10 -nnodes 1 -P $OLCF_PROJECT_ID -Is $SHELL
```

```
Job <993238> is submitted to default queue <batch>.
```

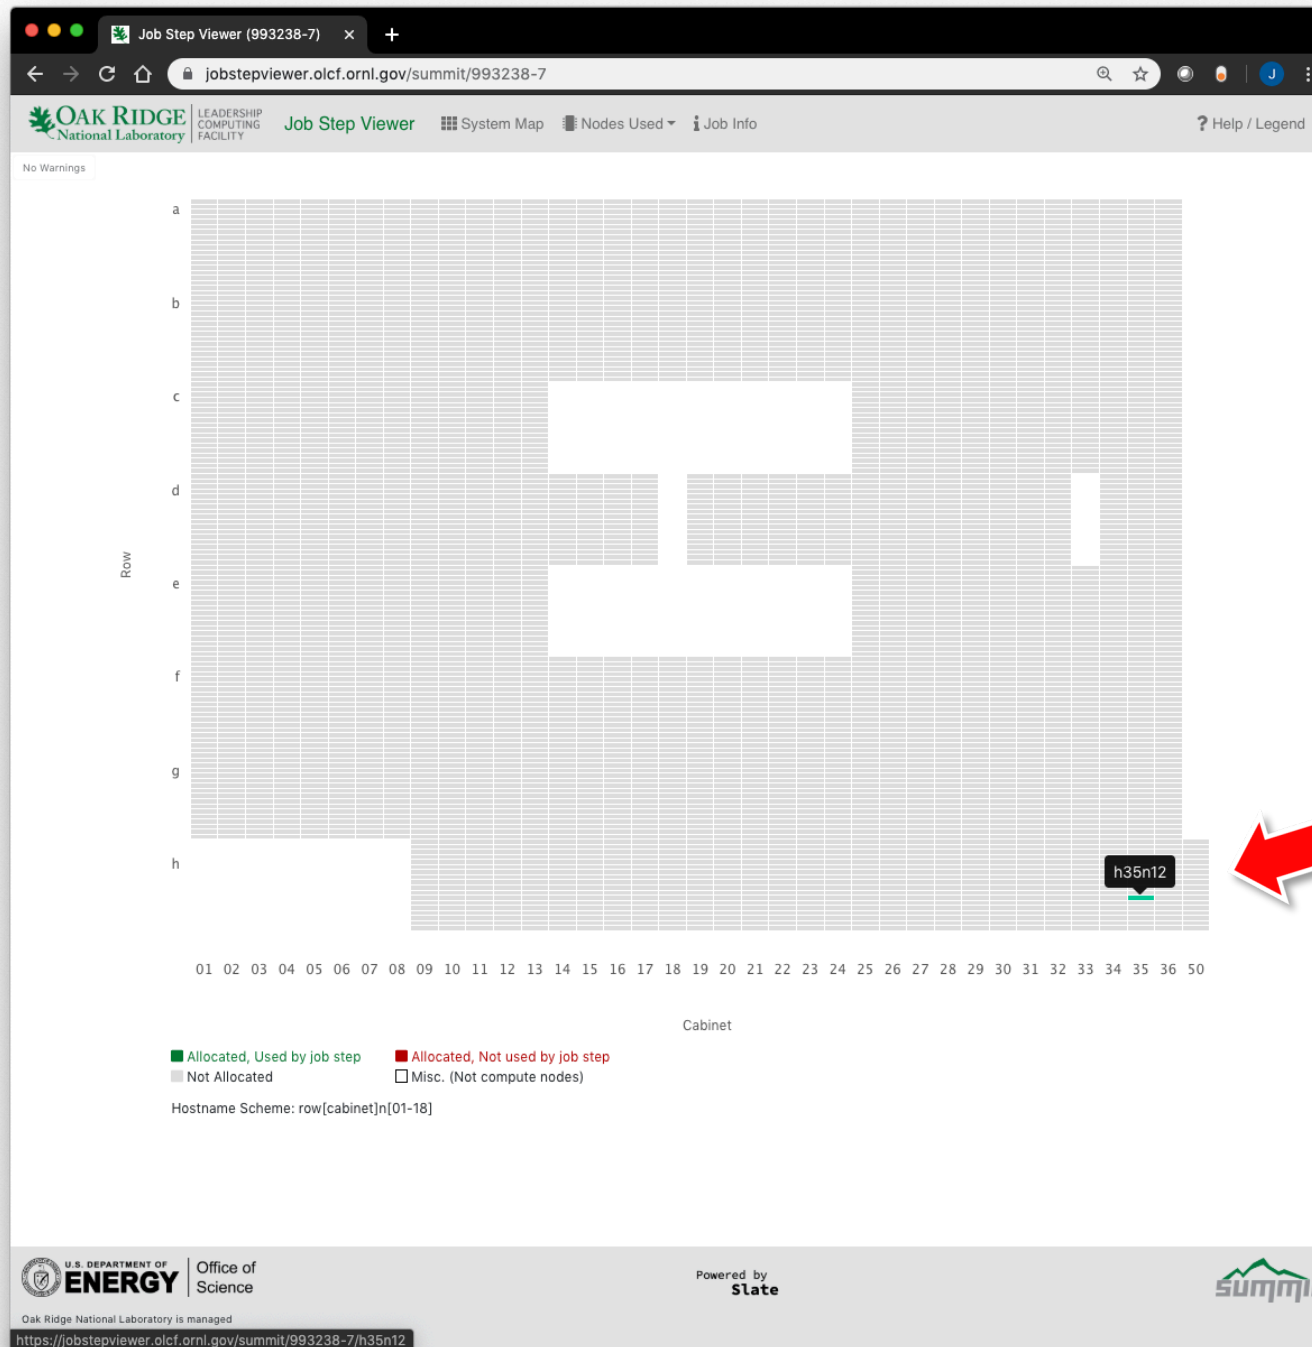
```
<<Waiting for dispatch ...>>
```

```
<<Starting on batch2>>
```

```
batch2.summit$ module load job-step-viewer
```

```
batch2.summit$ jsrun -n1 -a1 -c1 -g1
```

<https://jobstepviewer.olcf.ornl.gov/summit/993238-7>



Drill down into individual compute nodes

Job Step Viewer (993238-7)

jobstepviewer.olcf.ornl.gov/summit/993238-7/h35n12

OAK RIDGE  
National Laboratory

LEADERSHIP  
COMPUTING  
FACILITY

Job Step Viewer

System Map

Nodes Used

Job Info

Help / Legend

No Warnings

Summit Compute Node (h35n12)

IBM Power System AC922

jsrun -n1 -a1 -c1 -g1

Socket 0

0

1

2

0

001

028

029

056

057

002

003

030

031

058

059

004

005

032

033

060

061

006

007

034

035

062

063

008

009

036

037

064

065

010

011

038

039

066

067

012

013

040

041

068

069

014

015

042

043

070

071

016

017

044

045

072

073

018

019

046

047

074

075

020

021

048

049

076

077

022

023

050

051

078

079

024

025

052

053

080

081

026

027

054

055

082

083

Memory (256GB)

X-bus

Socket 1

3

4

5

088

089

116

117

144

145

090

091

118

119

146

147

092

093

120

121

148

149

094

095

122

123

150

151

096

097

124

125

152

153

098

099

126

127

154

155

100

101

128

129

156

157

102

103

130

131

158

159

104

105

132

133

160

161

106

107

134

135

162

163

108

109

136

137

164

165

110

111

138

139

166

167

112

113

140

141

168

169

114

115

142

143

170

171

Memory (256GB)

disabled

NVMe (1.6TB)

U.S. DEPARTMENT OF  
ENERGY

Office of  
Science

Powered by  
Slate

summit


Oak Ridge National Laboratory is managed  
by UT-Battelle for the US Department of Energy.

OAK RIDGE  
National Laboratory

9

# Summit Compute Node

- IBM Power System AC922 Compute Node ([Redbook](#))



# Features

- Easy to use
- Stand-alone jsrun testing
- Informed by real runs on Summit & respects user's runtime environment
- Warnings for common mistakes
- SMT{1,2,4} support
- NVMe (Burst Buffer) support
- Supports Explicit Resource Files (ERF)
- Easily sharable, unique URLs

```
$ jsrun -n6 -a1 -c7 -g1 -E OMP_NUM_THREADS=7 -brs
```

<https://jobstepviewer.olcf.ornl.gov/submit/993238-8>

OR

```
$ export OMP_NUM_THREADS=7
```

```
$ jsrun -n6 -a1 -c7 -g1 -brs
```



Job Step Viewer (993238-8)

+

jobstepviewer.olcf.ornl.gov/summit/993238-8/h35n12

🔍 ☆ 🌐 🇺🇸 J ⋮

OAK RIDGE  
National Laboratory

LEADERSHIP  
COMPUTING  
FACILITY

Job Step Viewer

System Map

Nodes Used ▾

Job Info

Help / Legend

No Warnings

### Summit Compute Node (h35n12)

IBM Power System AC922

```
jsrun -n6 -a1 -c7 -g1 -E <ENV_VAR> -brs
```

Socket 0

0

1

2

Memory (256GB)

0[0]	001	028	1[0]	2[0]	057
002	003	030	031	058	059
0[1]	005	1[1]	033	2[1]	061
006	007	034	035	062	063
0[2]	009	1[2]	037	2[2]	065
010	011	038	039	066	067
0[3]	013	1[3]	041	2[3]	069
014	015	042	043	070	071
0[4]	017	1[4]	045	2[4]	073
018	019	046	047	074	075
0[5]	021	1[5]	049	2[5]	077
022	023	050	051	078	079
0[6]	025	1[6]	053	2[6]	081
026	027	054	055	082	083

X-bus

Socket 1

3

4

5

Memory (256GB)

088	3[0]	116	117	144	145
090	091	118	4[0]	5[0]	147
3[1]	093	4[1]	121	5[1]	149
094	095	122	123	150	151
3[2]	097	4[2]	125	5[2]	153
098	099	126	127	154	155
3[3]	101	4[3]	129	5[3]	157
102	103	130	131	158	159
3[4]	105	4[4]	133	5[4]	161
106	107	134	135	162	163
3[5]	109	4[5]	137	5[5]	165
110	111	138	139	166	167
3[6]	113	4[6]	141	5[6]	169
114	115	142	143	170	171

disabled

NVMe (1.6TB)

U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science

Oak Ridge National Laboratory is managed  
by UT-Battelle for the US Department of Energy.

Powered by  
**Slate**

summit

OAK RIDGE  
National Laboratory

13

Job Step Viewer (993238-8)

jobstepviewer.olcf.ornl.gov/summit/993238-8/h35n12

OAK RIDGE

National Laboratory

LEADERSHIP

COMPUTING

FACILITY

Job Step Viewer

System Map

Nodes Used

Job Info

Help / Legend

No Warnings

Summit Compute Node (h35n12)

IBM Power System AC922

jsrun -n6 -a1 -c7 -g1 -E <ENV\_VAR> -brs

Socket 0

Socket 1

Memory (256GB)

Memory (256GB)

disabled

NVMe (1.6TB)

Process Notation	Meaning	Actions
#	Rank ID	N/A
#[#]	Rank ID [ Thread ID ]	N/A
...	Oversubscribed Hardware Thread (multiple ranks and/or multiple threads)	Hover to expand
+	Rank ID ([ Thread ID ]) too long to display	Hover to expand

0[6] 025 1[6] 053 2[6] 081 3[6] 113 4[6] 141 5[6] 169

026 027 054 055 082 083 114 115 142 143 170 171

U.S. DEPARTMENT OF


ENERGY

Office of Science

Oak Ridge National Laboratory is managed by UT-Battelle for the US Department of Energy.

Powered by Slate

summit

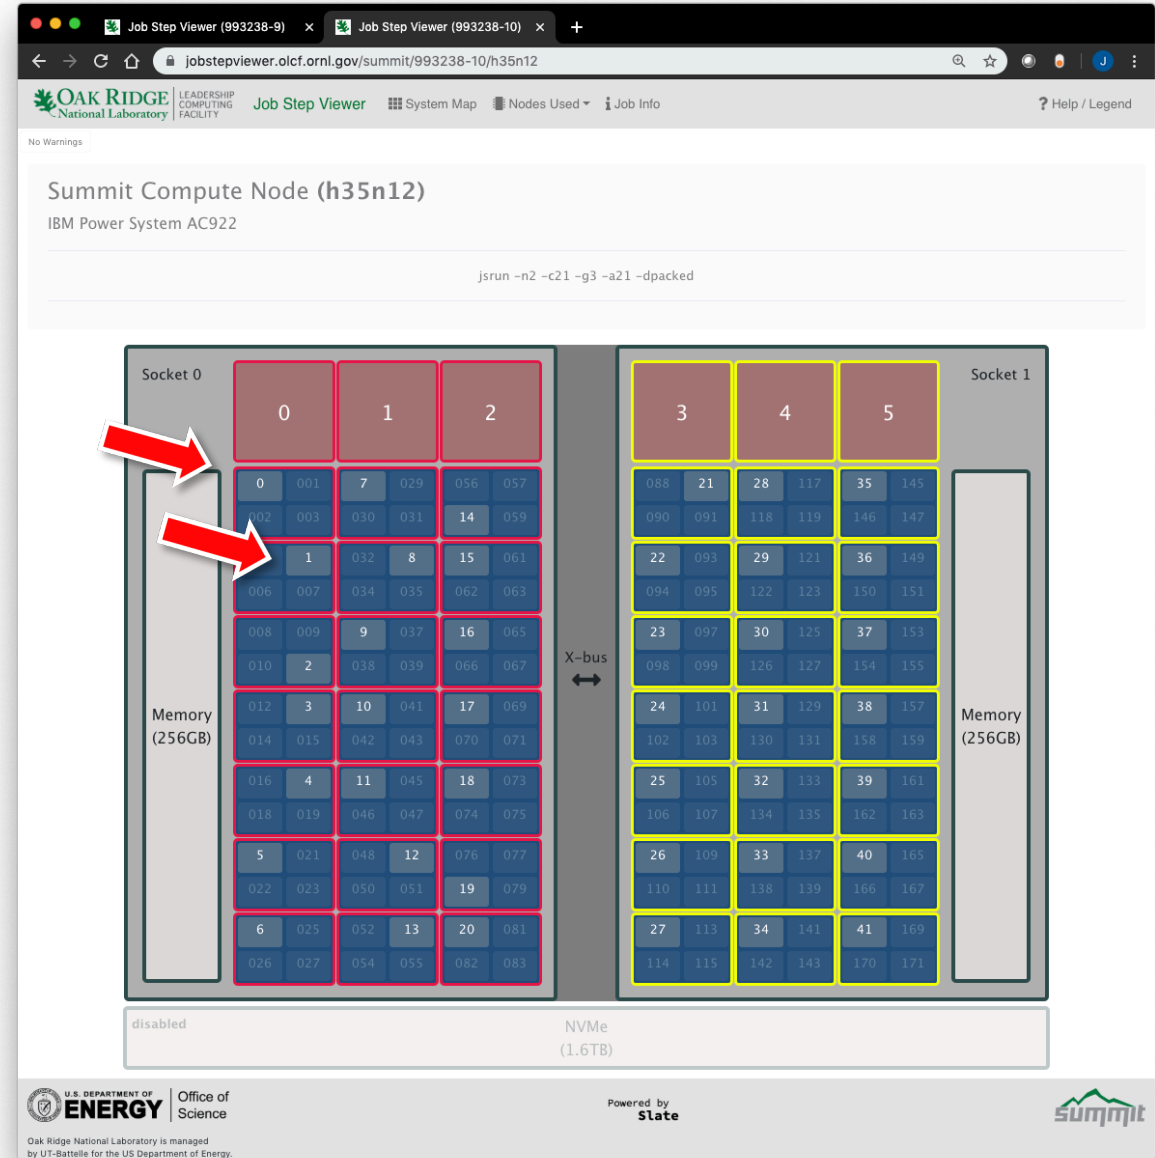
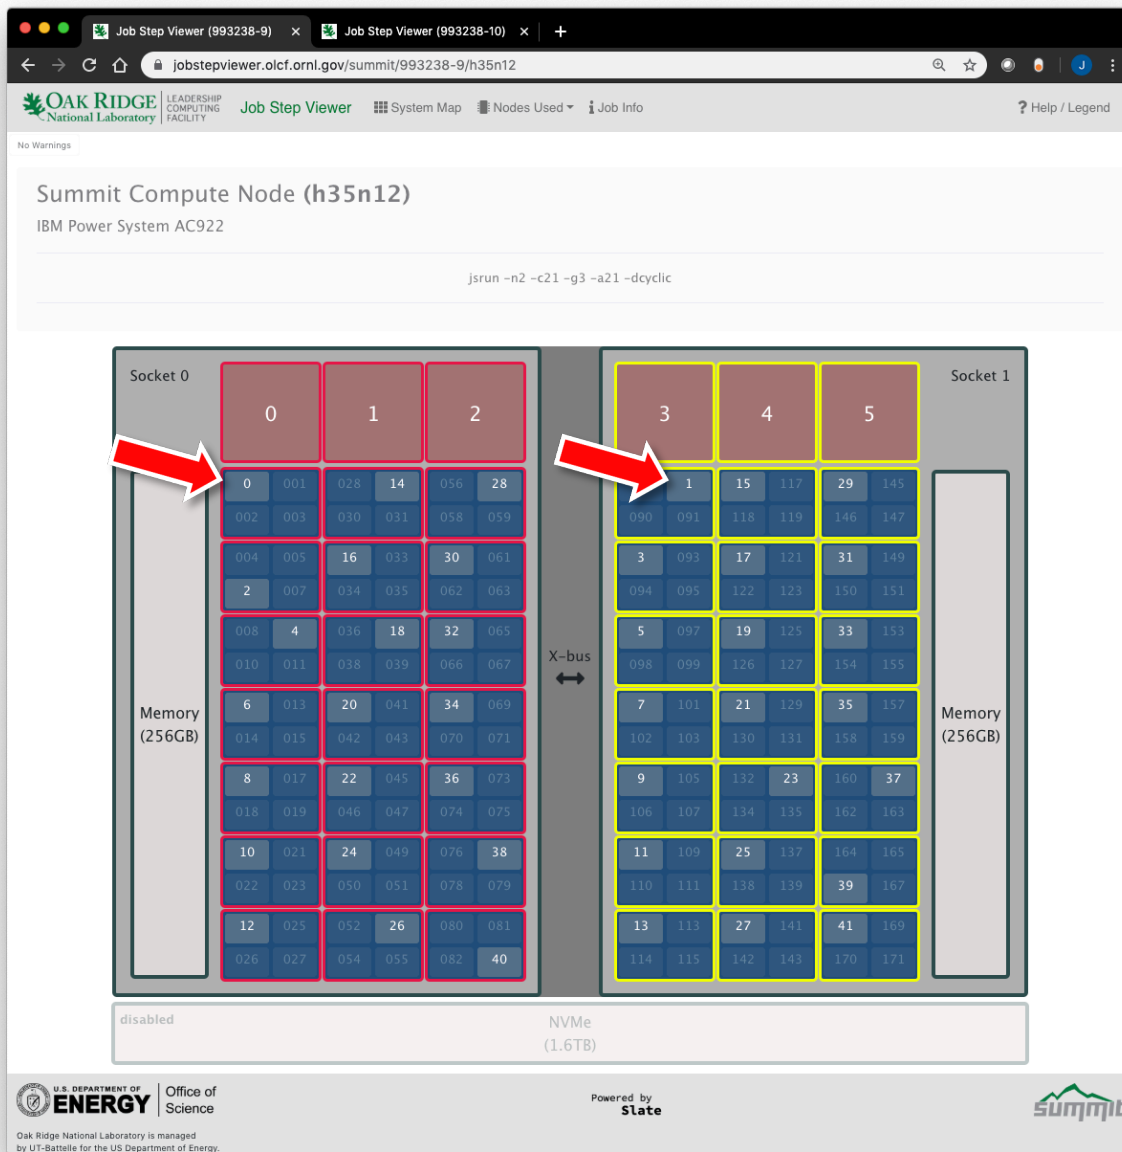


```
$ jsrun -n2 -c21 -g3 -a21 -dcyclic
```

<https://jobstepviewer.olcf.ornl.gov/submit/993238-9>

```
$ jsrun -n2 -c21 -g3 -a21 -dpacked
```

<https://jobstepviewer.olcf.ornl.gov/submit/993238-10>



# Features

- Easy to use
- Stand-alone jsrun testing
- Informed by real runs on Summit & respects user's runtime environment
- **Warnings for common mistakes**
- SMT{1,2,4} support
- NVMe (Burst Buffer) support
- Supports Explicit Resource Files (ERF)
- Easily sharable, unique URLs

```
$ jsrun -n30 -r6 -c7 -g1 -a1 -EOMP_NUM_THREADS=7
```

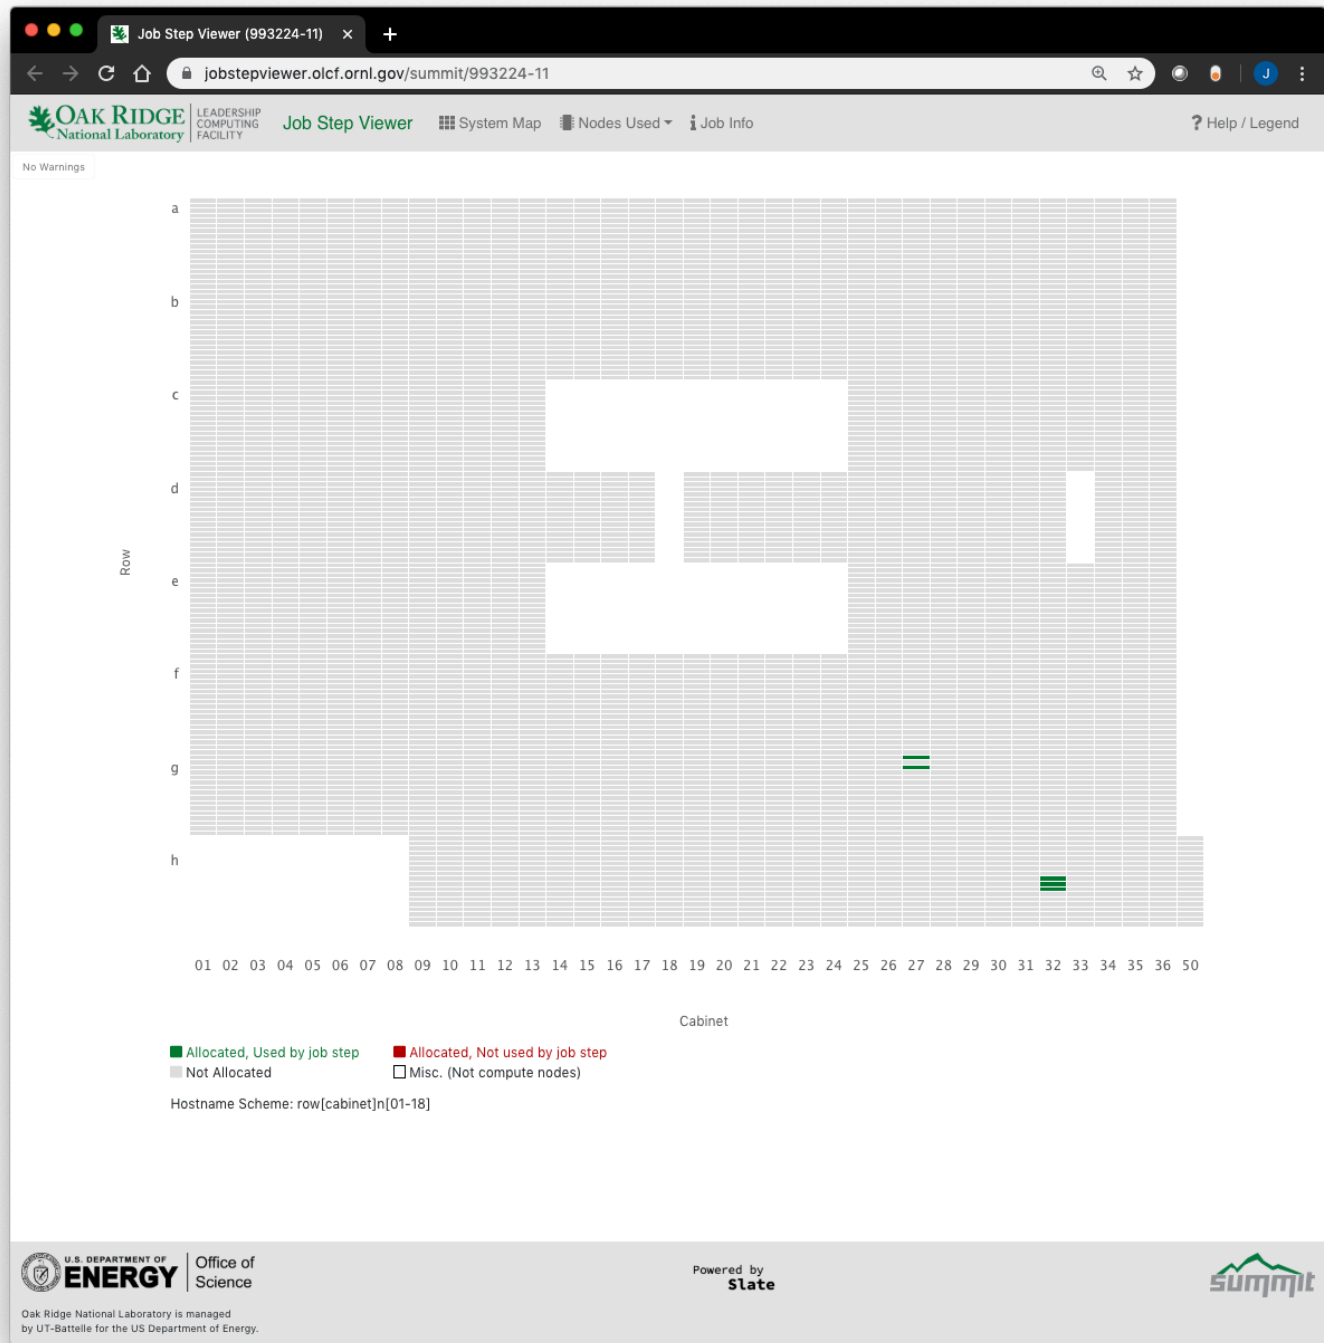
```
Warning: OMP_NUM_THREADS=7 is greater than available PU's
```

```
Warning: OMP_NUM_THREADS=7 is greater than available PU's
```

```
Warning: OMP_NUM_THREADS=7 is greater than available PU's
```

```
<snip>
```

<https://jobstepviewer.olcf.ornl.gov/summit/993224-11>





Job Step Viewer (993224-11)

jobstepviewer.olcf.ornl.gov/summit/993224-11/h32n10

OAK RIDGE

National Laboratory

LEADERSHIP

COMPUTING

FACILITY

Job Step Viewer

System Map

Nodes Used

Job Info

Help / Legend

Show/Hide Warnings

Warning: Oversubscribed Hardware Threads

Hardware Thread oversubscription may lead to decreased performance. Check task and thread binding.

### Summit Compute Node (h32n10)

IBM Power System AC922

```
jsrun -n30 -r6 -c7 -g1 -a1 -E <ENV_VAR>
```

Socket 0

0

1

2

...

...

19[0]

...

...

20[0]

18[6]

18[0]

...

19[6]

...

...

004

005

032

033

060

061

006

007

034

035

062

063

008

009

036

037

064

065

010

011

038

039

066

067

012

013

040

041

068

069

014

015

042

043

070

071

016

017

044

045

072

073

018

019

046

047

074

075

020

021

048

049

076

077

022

023

050

051

078

079

024

025

052

053

080

081

026

027

054

055

082

083

Memory (256GB)

X-bus

Socket 1

3

4

5

21[0]

...

22[0]

...

...

23[1]

...

21[1]

...

22[1]

...

23[0]

092

093

120

121

148

149

094

095

122

123

150

151

096

097

124

125

152

153

098

099

126

127

154

155

100

101

128

129

156

157

102

103

130

131

158

159

104

105

132

133

160

161

106

107

134

135

162

163

108

109

136

137

164

165

110

111

138

139

166

167

112

113

140

141

168

169

114

115

142

143

170

171

Memory (256GB)

disabled

NVMe (1.6TB)

U.S. DEPARTMENT OF

ENERGY

Office of Science

Oak Ridge National Laboratory is managed by UT-Battelle for the US Department of Energy.

Powered by Slate

summit

```
$ jsrun -n30 -r6 -c7 -g1 -a1 -EOMP_NUM_THREADS=7
```

```
Warning: OMP_NUM_THREADS=7 is greater than available PU's
```

```
Warning: OMP_NUM_THREADS=7 is greater than available PU's
```

```
Warning: OMP_NUM_THREADS=7 is greater than available PU's
```

```
<snip>
```

<https://jobstepviewer.olcf.ornl.gov/summit/993224-11>

```
$ jsrun -n30 -r6 -c7 -g1 -a1 -EOMP_NUM_THREADS=7 -brs
```

<https://jobstepviewer.olcf.ornl.gov/summit/993224-16>

Job Step Viewer (993224-16)

jobstepviewer.olcf.ornl.gov/summit/993224-16/h32n10

OAK RIDGE  
National Laboratory

LEADERSHIP  
COMPUTING  
FACILITY

Job Step Viewer

System Map

Nodes Used

Job Info

Help / Legend

No Warnings

Summit Compute Node (h32n10)

IBM Power System AC922

jsrun -n30 -r6 -c7 -a1 -E <ENV\_VAR> -brs

Socket 0

0

1

2

Memory (256GB)

000	18[0]	028	19[0]	056	20[0]
002	003	030	031	058	059
18[1]	005	19[1]	033	20[1]	061
006	007	034	035	062	063
18[2]	009	19[2]	037	20[2]	065
010	011	038	039	066	067
18[3]	013	19[3]	041	20[3]	069
014	015	042	043	070	071
18[4]	017	19[4]	045	20[4]	073
018	019	046	047	074	075
18[5]	021	19[5]	049	20[5]	077
022	023	050	051	078	079
18[6]	025	19[6]	053	20[6]	081
026	027	054	055	082	083

X-bus

Socket 1

3

4

5

Memory (256GB)

21[0]	089	22[0]	117	23[0]	145
090	091	118	119	146	147
21[1]	093	22[1]	121	23[1]	149
094	095	122	123	150	151
21[2]	097	22[2]	125	23[2]	153
098	099	126	127	154	155
21[3]	101	22[3]	129	23[3]	157
102	103	130	131	158	159
21[4]	105	22[4]	133	23[4]	161
106	107	134	135	162	163
21[5]	109	22[5]	137	23[5]	165
110	111	138	139	166	167
21[6]	113	22[6]	141	23[6]	169
114	115	142	143	170	171

disabled

NVMe (1.6TB)

U.S. DEPARTMENT OF  
ENERGY

Office of  
Science

Powered by  
Slate

summit

Oak Ridge National Laboratory is managed  
by UT-Battelle for the US Department of Energy.

OAK RIDGE  
National Laboratory

23

# Features

- Easy to use
- Stand-alone jsrun testing
- Informed by real runs on Summit & respects user's runtime environment
- Warnings for common mistakes
- SMT{1,2,4} support
- NVMe (Burst Buffer) support
- Supports Explicit Resource Files (ERF)
- Easily sharable, unique URLs

Job Step Viewer (993260-7)

jobstepviewer.olcf.ornl.gov/summit/993260-7/g27n08

OAK RIDGE

National Laboratory

LEADERSHIP

COMPUTING

FACILITY

Job Step Viewer

System Map

Nodes Used

Job Info

Help / Legend

No Warnings

Summit Compute Node (g27n08)

IBM Power System AC922

jsrun -n6 -c7 -g1 -a14 -dcyclic

Socket 0

012

042143244

486497850

125413555614

186019612062

662425672668

723031737432

367879378038

Memory (256GB)

X-bus

Socket 1

345

453464475

51952101153

571558165917

216364222365

276928702971

753376343577

398140828341

Memory (256GB)

disabled

NVMe (1.6TB)

U.S. DEPARTMENT OF

ENERGY

Office of Science

Powered by

Slate

summit

Oak Ridge National Laboratory is managed by UT-Battelle for the US Department of Energy.

OAK RIDGE  
National Laboratory

25

# Features

- Easy to use
- Stand-alone jsrun testing
- Informed by real runs on Summit & respects user's runtime environment
- Warnings for common mistakes
- SMT{1,2,4} support
- NVMe (Burst Buffer) support
- Supports Explicit Resource Files (ERF)
- Easily sharable, unique URLs

```
$ cat erf.demo
```

```
cpu_index_using: physical
```

```
1 : {host: 1; cpu: {0:28};   gpu: 0}
```

```
1 : {host: 1; cpu: {28:28};  gpu: 1}
```

```
1 : {host: 1; cpu: {56:28};  gpu: 2}
```

```
1 : {host: 1; cpu: {88:28};  gpu: 3}
```

```
1 : {host: 1; cpu: {116:28}; gpu: 4}
```

```
1 : {host: 1; cpu: {144:28}; gpu: 5}
```

```
$ jsrun --erf_input $HOME/erf.demo
```

<https://jobstepviewer.olcf.ornl.gov/summit/993224-3>

Job Step Viewer (993224-3)

+

jobstepviewer.olcf.ornl.gov/summit/993224-3/g27n03

OAK RIDGE

National Laboratory

LEADERSHIP

COMPUTING

FACILITY

Job Step Viewer

System Map

Nodes Used

Job Info

Help / Legend

No Warnings

Summit Compute Node (g27n03)

IBM Power System AC922

jsrun --erf\_input <PATH>

Socket 0

0

1

2

Memory (256GB)

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027

028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 054 055

056 057 058 059 060 061 062 063 064 065 066 067 068 069 070 071 072 073 074 075 076 077 078 079 080 081 082 083

Socket 1

3

4

5

Memory (256GB)

088 089 090 091 092 093 094 095 096 097 098 099 100 101 102 103 104 105 106 107 108 109 110 111 112 113 114 115

116 117 118 119 120 121 122 123 124 125 126 127 128 129 130 131 132 133 134 135 136 137 138 139 140 141 142 143

144 145 146 147 148 149 150 151 152 153 154 155 156 157 158 159 160 161 162 163 164 165 166 167 168 169 170 171

disabled

NVMe (1.6TB)

U.S. DEPARTMENT OF

ENERGY


Office of Science

Oak Ridge National Laboratory is managed by UT-Battelle for the US Department of Energy.

Powered by

Slate

summit


**OAK RIDGE**  
 National Laboratory

28



# Not Currently Supported

- MPMD modes
- `jsrun` *Process Management* options
  - `-e / --stdio_mode`
  - `-I / --stdin_rank`
  - `-k / --stdio_stderr`
  - `-o / --stdio_stdout`
  - `-t / --stdio_input`
- Other IBM Power Systems (Ascent, SummitDev,...)
- Profilers (`nvprov`)

If you feel we're missing a feature, please let us know!

# Powered by Slate

- Job Step Viewer is built and runs with normal user permissions
- What is Slate?
  - User-facing service designed for hosting long running processes such as workflows, data portals, and databases
- Get in touch with us if you have a use case for such a service!

# Job Step Viewer

A new utility for working with  
`jsrun` on Summit

Jack Morrison  
HPC Engineer, User Assistance  
March 2020

ORNL is managed by UT-Battelle, LLC for the US Department of Energy



U.S. DEPARTMENT OF  
**ENERGY**