

# *Resource Management for Next-generation HPC Systems: Challenges and Solutions*

*11<sup>th</sup> Scheduling for Large-scale Systems Workshop*

Tapasya Patki

May 20, 2016



# LLNL HPC Systems

System (Program)	Processor Architecture	Nodes	Cores	Peak (TFLOP/s)
RZ				
<a href="#">RZCereal</a> (M&IC)	Intel Xeon E5530	21	169	1.6
<a href="#">RZHasGPU</a>	Intel Xeon E5-2667 v3	20	320	8.2
<a href="#">RZMerl</a> (ASC/M&IC)	Intel Xeon E5-2670	162	2,592	53.9
<a href="#">RZSLIC</a> ***	Intel Xeon E5330	3		
<a href="#">RZuSeq</a> (ASC) ****	IBM PowerPC A2	522		
<a href="#">RZzeus</a> (M&IC)	Intel Xeon E5530	267		

System (Program)	Processor Architecture	Nodes	Cores	Peak (TFLOP/s)
CZ				
<a href="#">Ansel</a> (M&IC)	Intel Xeon EP X5660	324	3,888	43.5
<a href="#">Aztec</a> (M&IC)	Intel Xeon EP X5660	96	1,152	12.9
<a href="#">Catalyst</a> (ASC/M&IC) ****	Intel Xeon E5-2695 v2	324	7,776	149.3
<a href="#">Cab</a> (ASC/M&IC)	Intel Xeon E5-2670	1,296	20,736	431.3

**Stats**

Max: 98,304 nodes in one system (Sequoia)

25 systems across open and closed zones

Various processor architectures

<a href="#">Inca</a> (ASC)	Intel X5660						
<a href="#">Juno</a> (ASC)	AMD 8354						
<a href="#">Max</a> (ASC)	Intel Xeon E5-2670						
<a href="#">Muir</a> (ASC)	Intel Xeon EP X5660	1,296	15,552	174.2			
<a href="#">Sequoia</a> (ASC) **	IBM PowerPC A2	98,304	1,572,864	20,132			
<a href="#">Zin</a> (ASC)	Intel Xeon E5-2670	2,916	46,656	970.4			

# Future systems present several new challenges

---

- Run-to-run variability, inter-job interference
- **Multiple constraints**: power, network, I/O, and data awareness
- Need to support high-throughput workloads, such as UQ workloads
- Increased error and failure rates

# Future systems present several new challenges

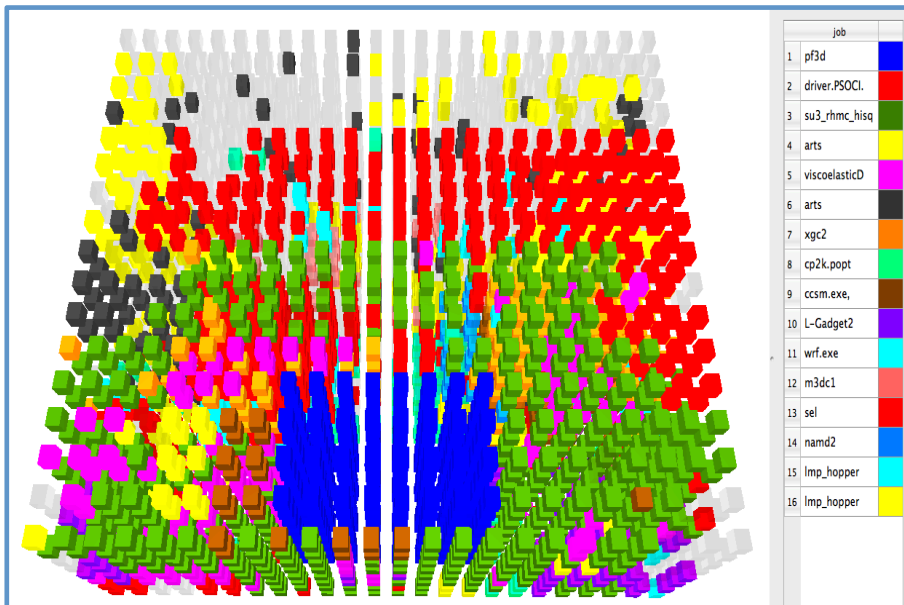
- Run-to-run variability, inter-job interference
- **Multiple constraints**: power, network, I/O, and data awareness
- Need to support high-throughput workloads, such as UQ workloads
- Increased error and failure rates

*How do we design low-overhead, scalable resource managers?*


# **PROBLEM:** Network contention and inter-job interference can lead to severe run-to-run variability

Two 512-node pF3D runs (blue) on Hopper Cray XE6

*Slow*



*Fast*



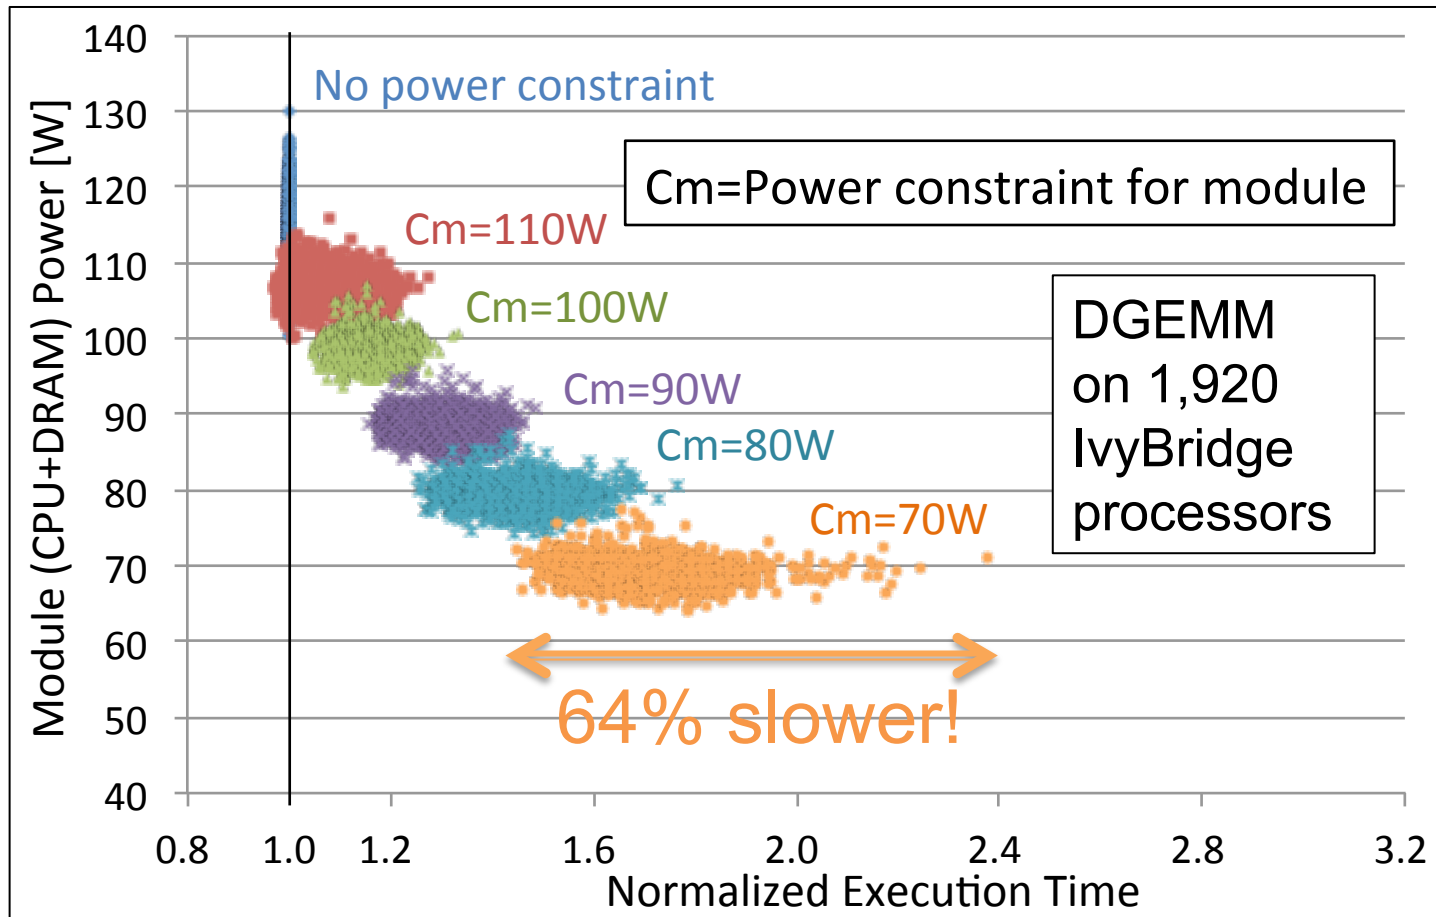
**25% slower messaging rate due to other jobs!**

*Graph Courtesy: Abhinav Bhatele, LLNL*

# **PROBLEM:** Existing resource managers may not scale well for high-throughput ensemble workloads

- Managing several thousand jobs at once can be slow, e.g., UQ ensemble workloads
  - Job launch times vary
- Typically targeted toward fewer large-sized jobs, but not for more, small-sized jobs
- ~250 jobs on clusters that have thousands of nodes

# **PROBLEM:** Power capping affects each processor differently, creating runtime variability and imbalance



Inadomi et al., SC'15

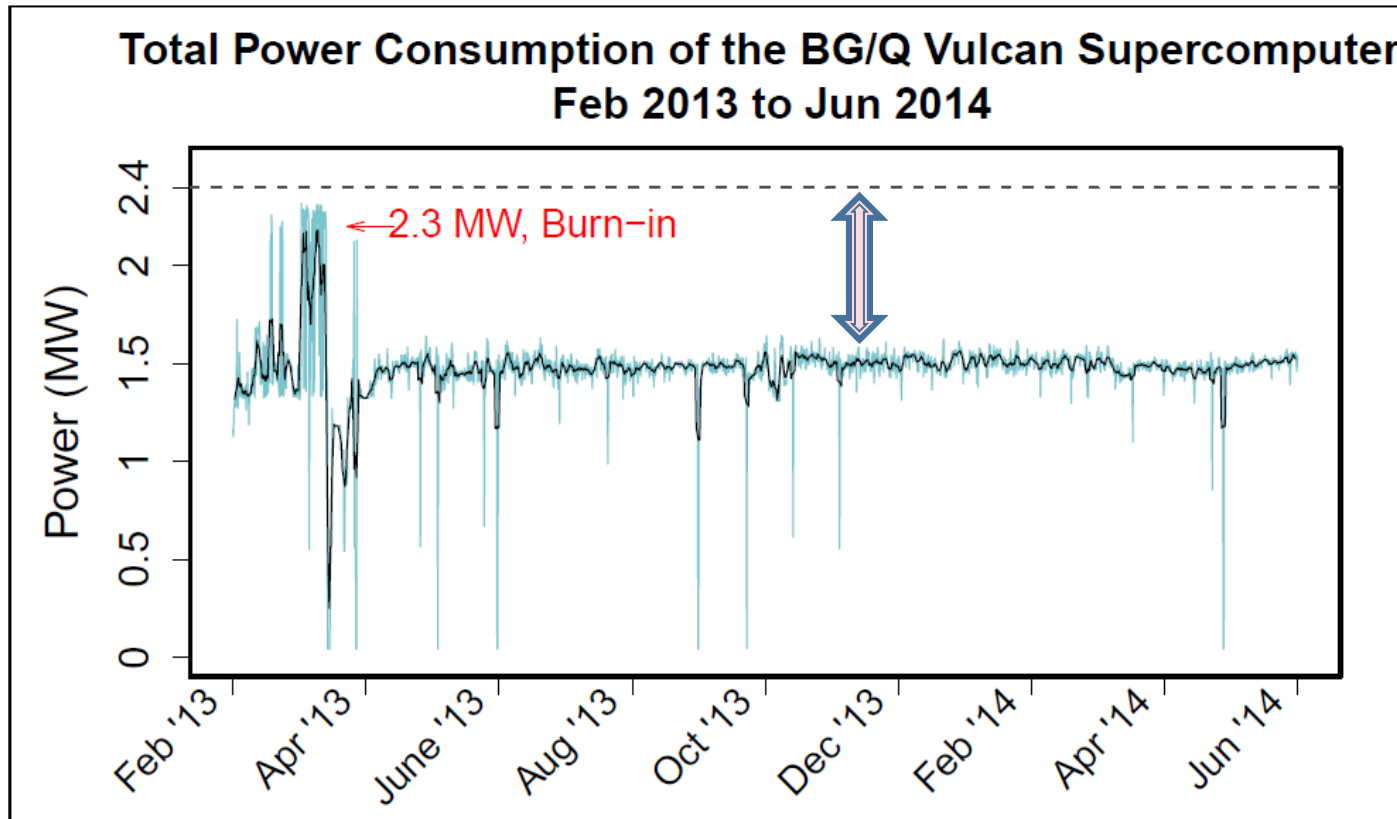
# ***RESULT: Unhappy Users!***

---





# ***UNDERUTILIZED CAPACITY:*** Utilizing one resource well leads to under-utilization of another resource



**40%  
Procured  
Power  
Unused!**

*Patki et al., HPDC'15*

# Existing resource managers **cannot be extended** to support multi-constraint HPC systems easily

	Network Topology	Network Bandwidth	I/O-Aware	Power-Aware	Type?
SLURM	✓	✗	✗	⊙	Monolithic
Moab/ Torque	✓	✗	✗	✗	Monolithic
LSF, IBM	✓	✗	⊙	⊙	Monolithic
Cobalt, Argonne	✓	✗	✗	✗	Monolithic
Mesos, Apache	✓	✗	✗	⊙	2-Level Hierarchical
PBSPPro, Altair	✓	✗	✗	✗	Monolithic

⊙: Limited support,  
In progress

# Existing resource managers are **not** designed to be fault-tolerant

---

- Moldable and malleable jobs are typically not supported
- Checkpoint/Restart process is slow and un-intelligent
- Users end up requesting 'redundant' nodes and more time as part of allocation
  - Underutilized capacity (see Felix's talk!)

# ***RESULT: Unhappy System Administrators!***

---



# LLNL's approach is to provide a holistic solution for a large-scale HPC system

Monitoring  
infrastructure for  
production clusters

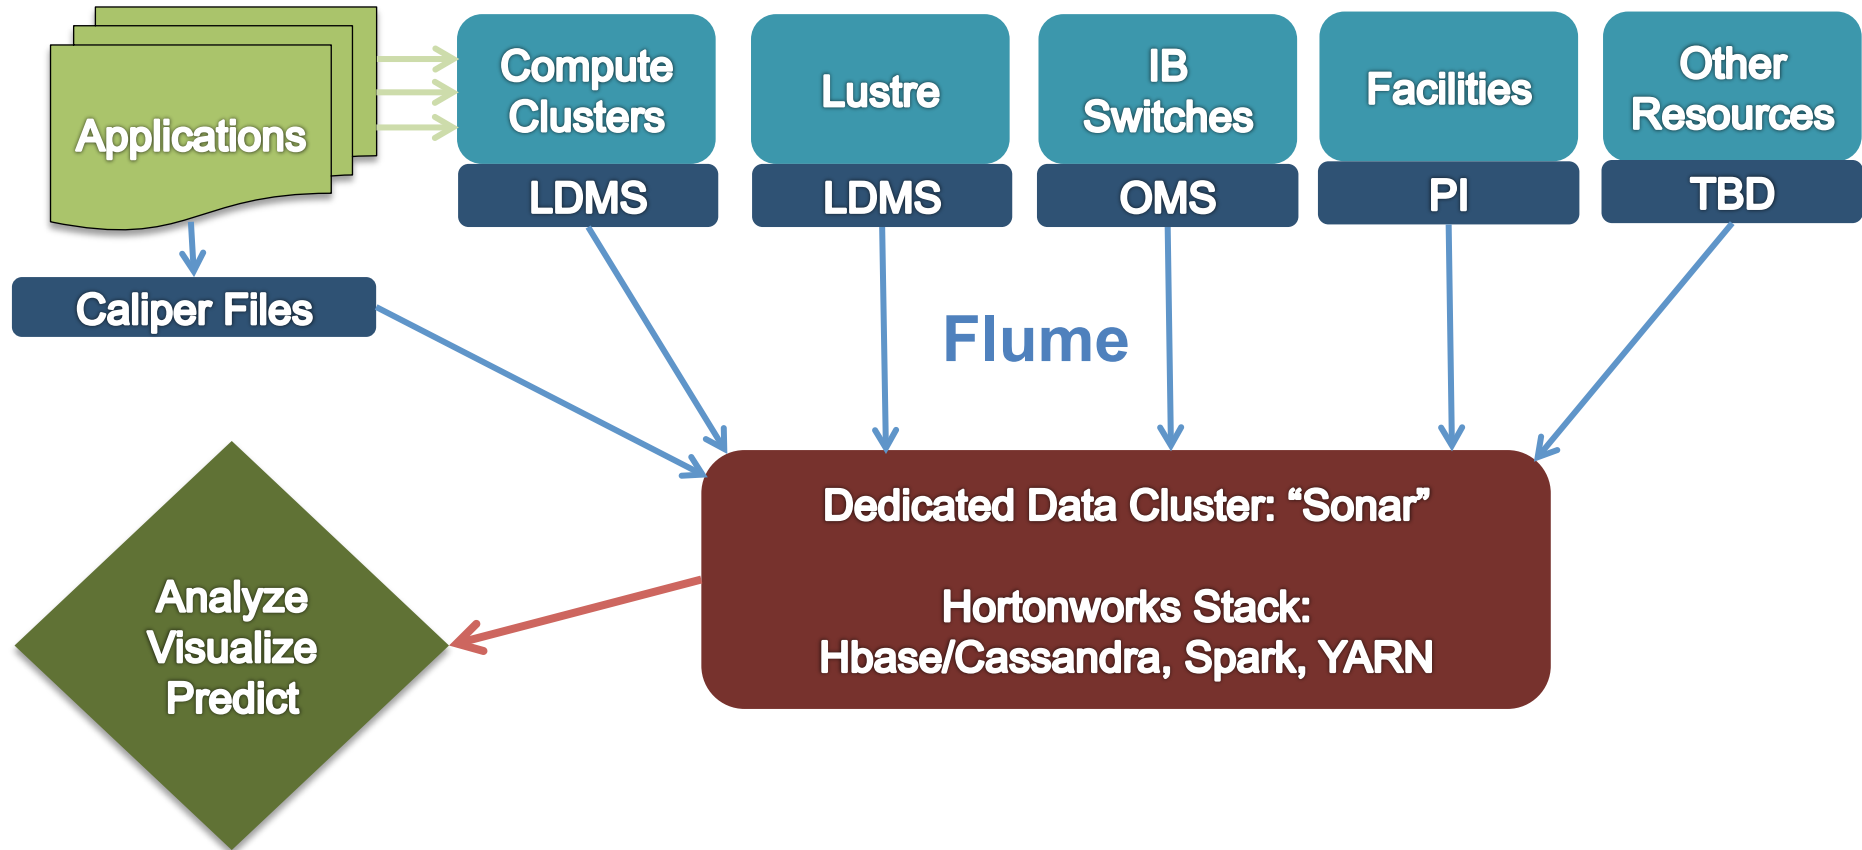
*PIPER\**, SONAR cluster

Framework for next-  
gen resource  
management



*\*Performance Insights for Programmers  
and Exascale Runtimes*

# Deploying site-wide monitoring infrastructure



PI: Todd Gamblin, LLNL

# Flux Framework: Next Generation Resource Mgmt.

---

- **Hierarchical** resource manager designed to support future HPC systems in a scalable manner
- Three key components: flux-core, flux-sched, flux-capacitor
- Open source, initial release will be available soon



# Flux Framework: *flux-core*

- Communication layer comprising of distributed message broker and plug-in modules for services
- Three overlay networks implemented using ZeroMQ
  - For request/response, session-wide broadcasts, and debugging
- Flux KVS/DHT module for job and resource configuration in a session
  - Useful for logging, synchronization, broadcasting



# Flux Framework: *flux-sched*

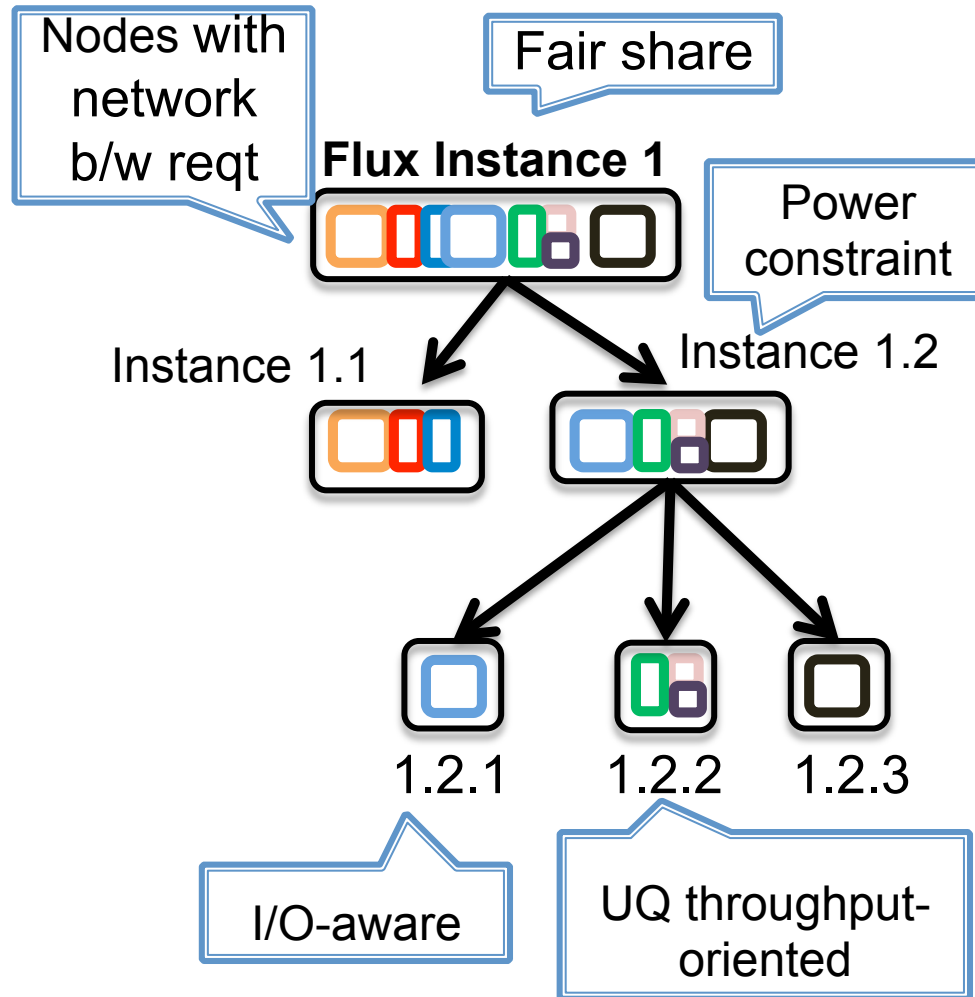
- Flexible resource model to represent new **hardware** and **flow** resources
  - Accelerators
  - Power, network bandwidth
  - Easily extensible and mutable
- Scalable job management with asynchronous event-based protocols
- Currently supports FCFS, easy backfilling

# Flux Framework: *flux-capacitor*

---

- Simple python interface for high-throughput, ensemble workloads
- Supports pulsed job launch: feeds jobs (flux instances) to the system at an ideal rate

# Flux Framework: Example



# Flux Framework: Current Research

- Hierarchical Scheduling
  - Adds additional levels of schedulers to form a hierarchy
  - Initial study shows adding only one additional level in the **reduces the scheduling complexity by 3.6x**
  - But increases the resource fragmentation by up to 20%
- Dynamic Scheduling
  - Allows for allocations to change size at runtime
  - Fault-tolerance, fragmentation, utilization
- I/O Aware Scheduling
  - Efficient ways to schedule for systems with burst buffers
  - Increases system efficiency by  $\sim 1.3x$  in exchange for increasing turnaround time by  $\sim 1.5x$

# Open Research Questions

---

- Sources of run-to-run variability
- Analysis of user behavior, workloads, error logs
- Impact of data staging on power and network performance
- How to prioritize constraints: power, network, file systems


# The Flux Team



Lawrence Livermore National Laboratory\*

University of Delaware

Technische Universität Darmstadt



\*And Mark Grondona, whose picture I couldn't find this morning!



Project Page: [flux-framework.github.io](https://flux-framework.github.io)

flux-core: [github.com/flux-framework/flux-core](https://github.com/flux-framework/flux-core)

flux-sched: [github.com/flux-framework/flux-sched](https://github.com/flux-framework/flux-sched)

flux-capacitor: [github.com/flux-framework/capacitor](https://github.com/flux-framework/capacitor)

