Cray XC System Software Update

The Beat Goes On

kmark@cray.com

https://www.linkedin.com/in/kelly-mark

AGENDA

- Introducing CLE 7.0!
  - Deep dive into Hugepage memory accounting
- XC Patching Strategy Change
- XC Software Roadmap
- XC Software Support Matrix
- Summary
- Q&A

WHY THE EXCITEMENT?

XC – IN THE NEWS

BRIDGE TO SHASTA

INTRODUCING CLE 7.0

Consistent upgrade process – same time requirements

- Same as CLE 6.0 UP06 -> CLE 6.0 UP07

SLES15, Lustre 2.11, NVIDIA 10.0

Performance improvements for ARM

- Turbo
- Native hugepages

Other Improvements

- SSA Client
- Logging & Networking
- eLogin/CLE common behavior (Ansible plays for standardized networking, ntp)
- Direct Attached Lustre (DAL) moved from CentOS to SLES
- Hugepage memory accounting

CLE 7.0 UP00 - Hugepages – Memory cgroup limits

- Memory cgroups now count and include hugepage usage in the overall memory usage
- Now the memory limits set by a WLM and enforced by the kernel will be more comprehensive
- Some jobs may fail because their use of hugepages is correctly tracked and counted against their memory cgroup limits
- This tracking is enabled by default but can be disabled with a boot time parameter

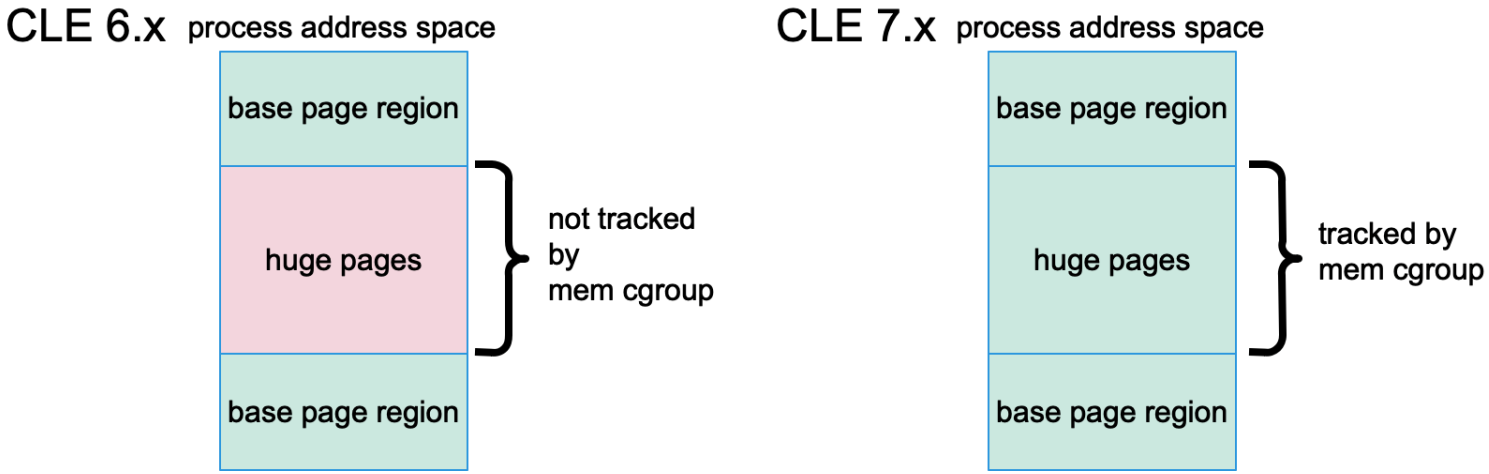

CLE 7.0 UP00 - Hugepages – Out of memory (OOM) score

- Reserved hugepages are now counted in addition to in use hugepages by the out-of-memory (OOM) killer when computing the OOM score for each process
- This means that the OOM killer will be better able to target the correct process
- This behavior can be disabled at runtime

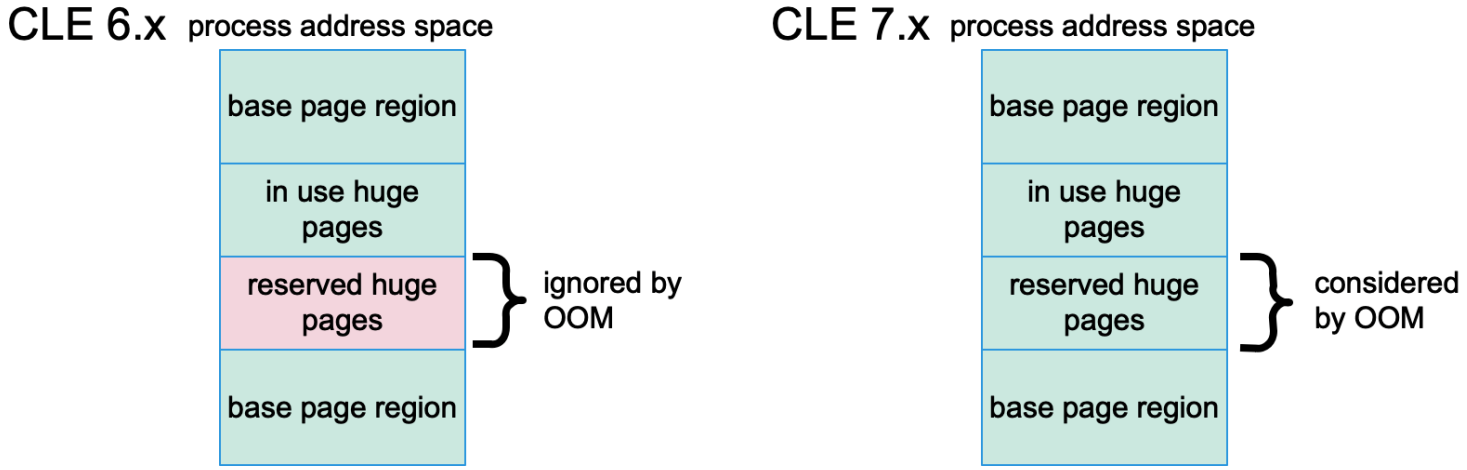

CLE 7.0 UP00 - Hugepage – Resident Set Size (RSS)

- The resident set size (RSS) calculations performed by the kernel now include all hugepages that are in use or reserved by the process
- As an example, the RSS data shown in /proc/<pid>/status, the summary line displayed by aprun after a job completes, and the max_rss value gathered by the RUR taskstats plugin will all now include hugepage usage
- This behavior can be disabled at runtime

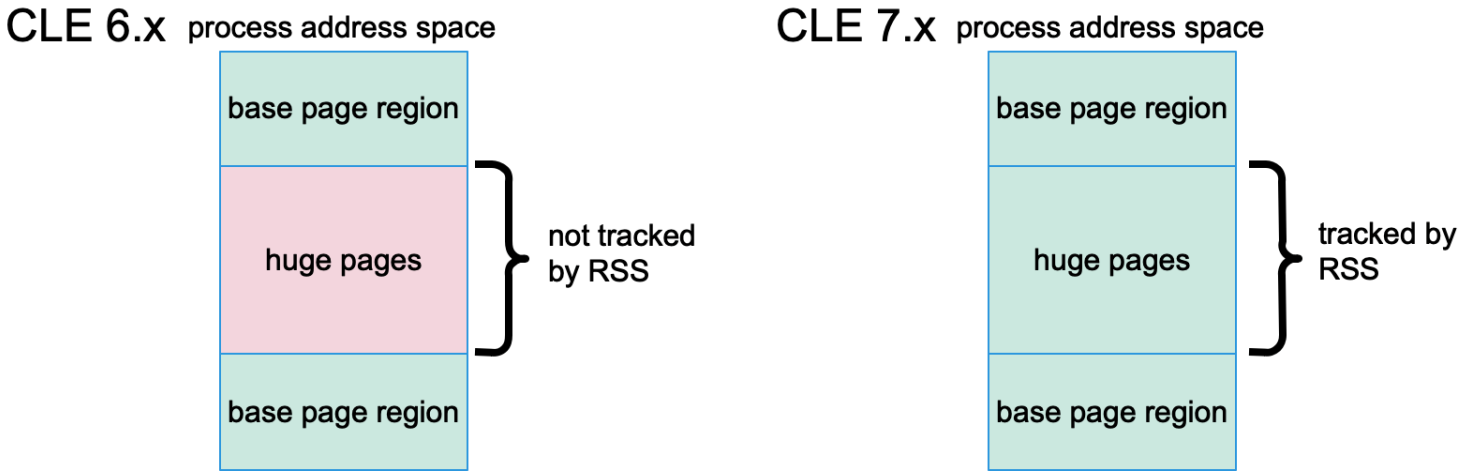

XC PATCH STRATEGY

Cumulative Patching

XC PATCH STRATEGY GOALS

- **Patches are Released as Generally Available**
  - Enables patches to be fully tested like a standard release
- **Standardize Installation Across all Patchsets**
  - Leverage tools used for Major/Minor releases
  - Reduces install errors
- **Eliminate Patch Dependency Chains**
  - Patches are cumulative – and contain any dependencies
- **Patches are Released on a Regular Cadence**
  - Monthly (or as needed)

Patches are generally available (GA) to all customers at the same time

WHAT IS IN A CUMULATIVE PATCHSET?

- README
  - One README.txt instead of three
  - Install instructions removed (unless exception)
- Standard installer (and removed INSTALL and LOAD scripts)
- Patchset manifest – rpm list of what changed in SMW/CLE ISOs
- CLE, SMW, SLE Update ISO's
- Revised Publication: S-2559
  - Standard instructions regardless of Patchset content

XC PATCH UPDATE - CURRENT

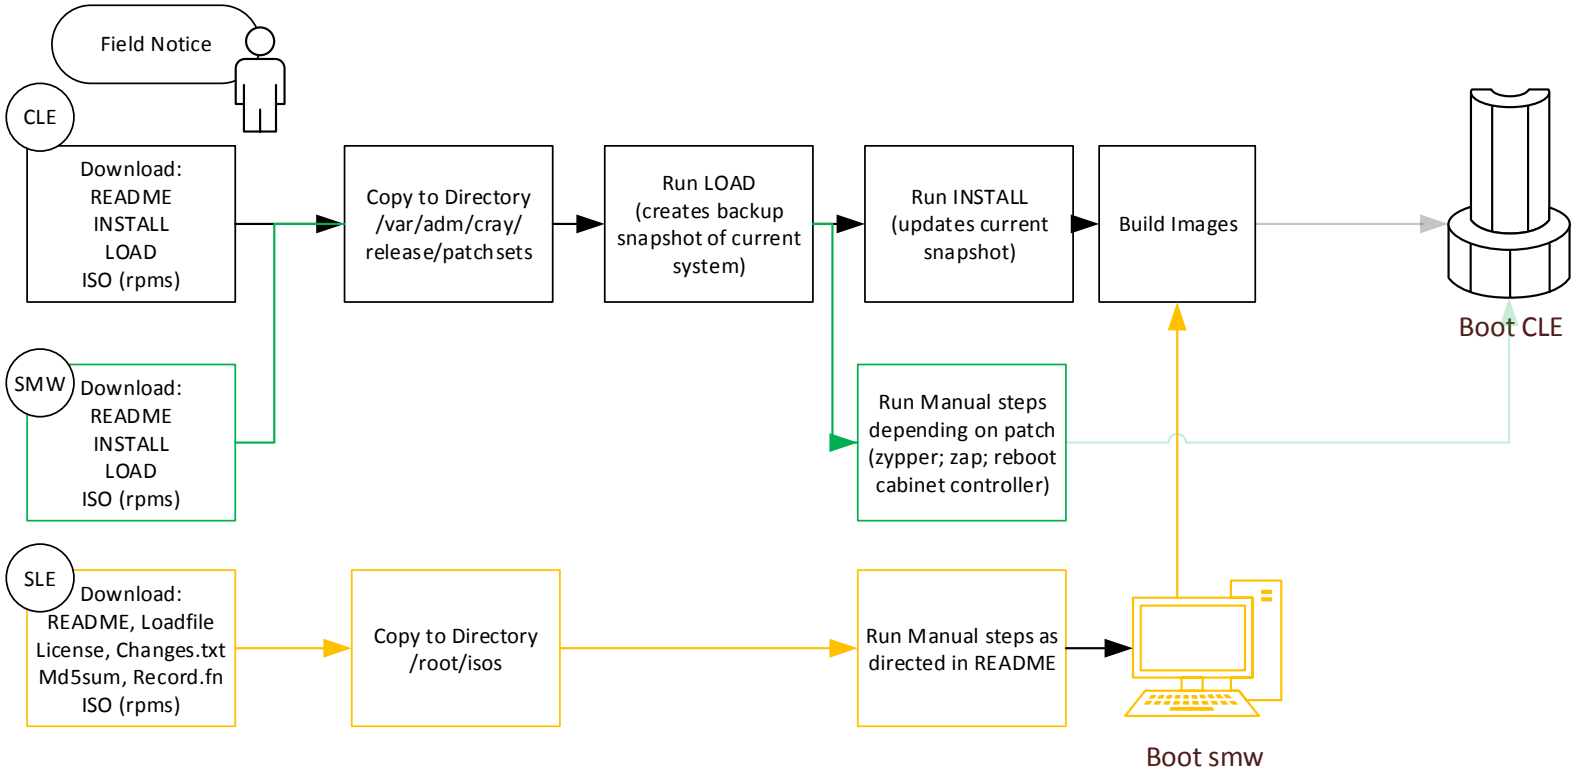

3 patch processes depending on patch type

Each patch required unique instructions

Inconsistent snapshot instructions

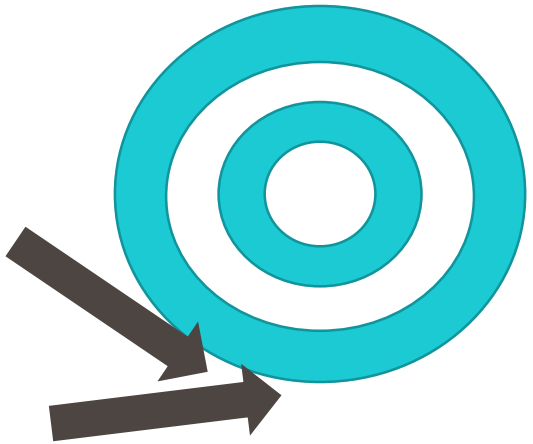

XC CUMULATIVE PATCH PROCESS

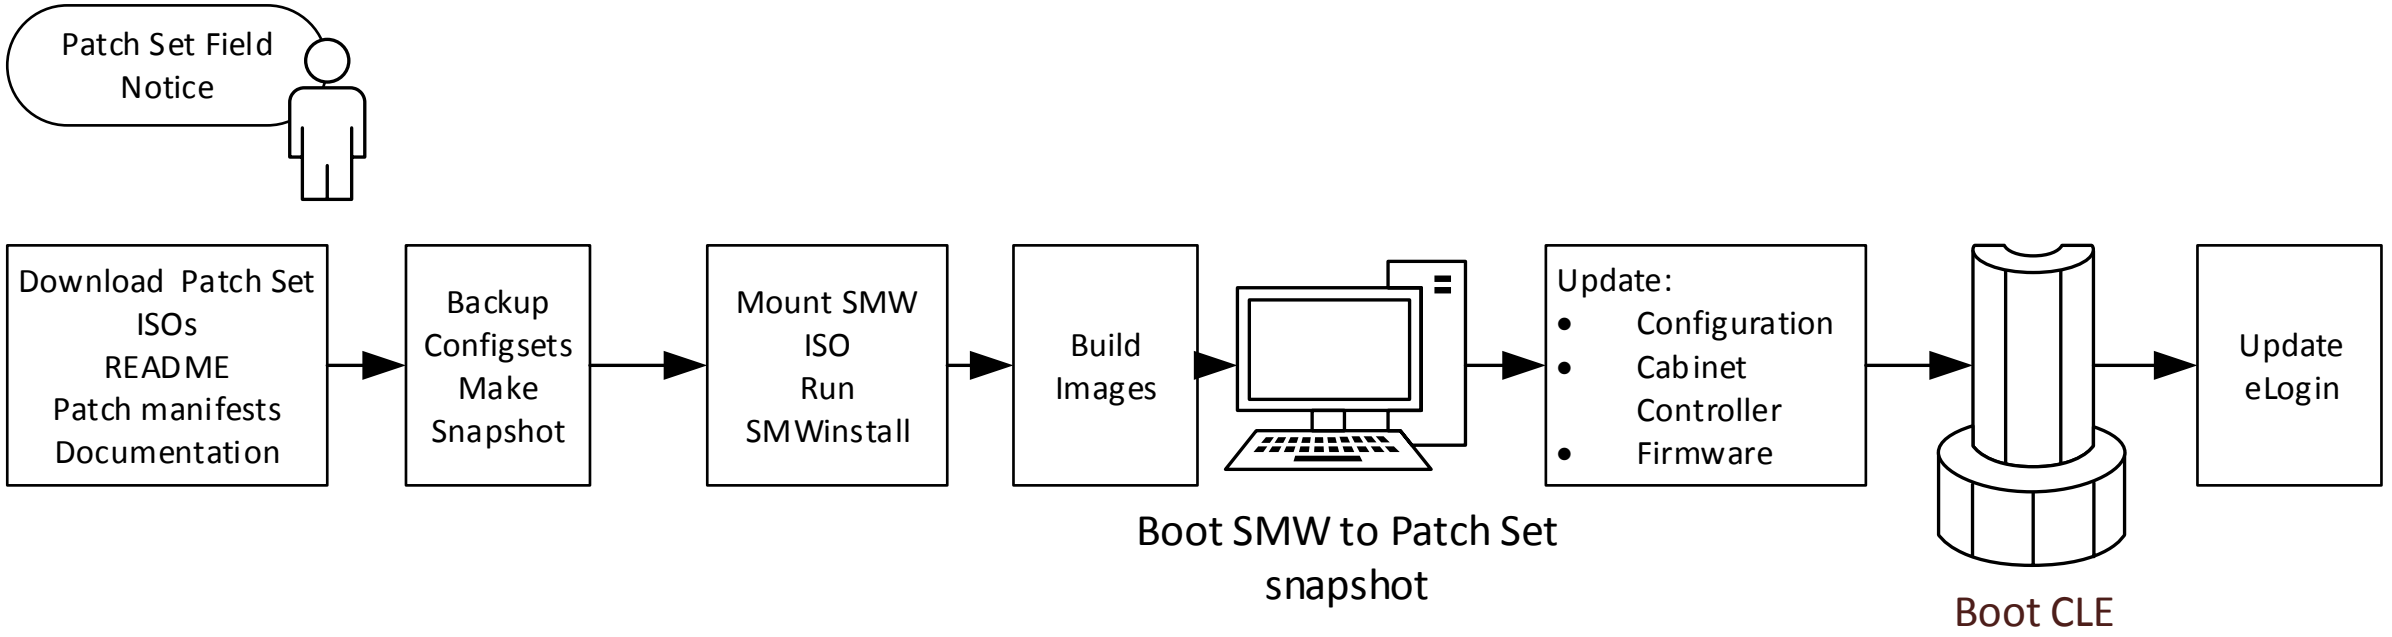

1 patch process

Same process as software updates

Fully documented, few exceptions

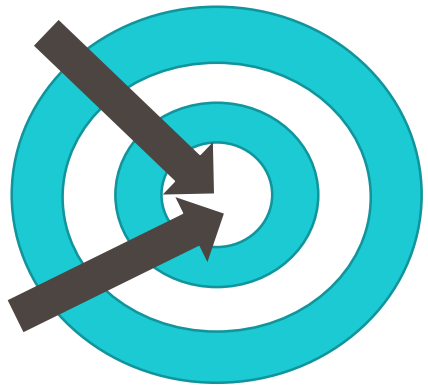

CUMULATIVE PATCHSET - EXAMPLE

	PS01 March	PS02 April	PS03 May
CLE 7.0.UP00	March changes	March changes	March Changes
		April Changes	April Changes
			May Changes
SMW 9.0.UP00	March Changes	March Changes	March Changes
		April Changes	April Changes
SLE Update		SLE Update-April	SLE Update April

XC PATCH UPDATE – SCHEDULE

Release	April	May	June	July	August
CLE 7.0/SMW 9.0 UP00	PS02 4/3	PS03 5/1	Week1	Week1	Week1
CLE 6.0/SMW 8.0 UP07	PS31 4/10	PS32 5/8	Week2	Week2	Week2
CLE 6.0/SMW 8.0 UP06	n/a	PS36 5/15	Week3	Week3	Week3

XC PATCH UPDATE – CUSTOMER FEEDBACK

- NO COMPLAINTS (yet)!!
  - 10 customers have applied 6.0.UP07 PS30
  - 6 customers have applied 7.0.UP00 PS02

- Some praised new approach

- A few questions
  - One customer was concerned about patch size, we were able to remove multiple CUDA rpms and significantly reduced ISO size
  - One customer noticed a documentation error in PS30

XC ROADMAP

Software

2019 XC ROADMAP – SOFTWARE

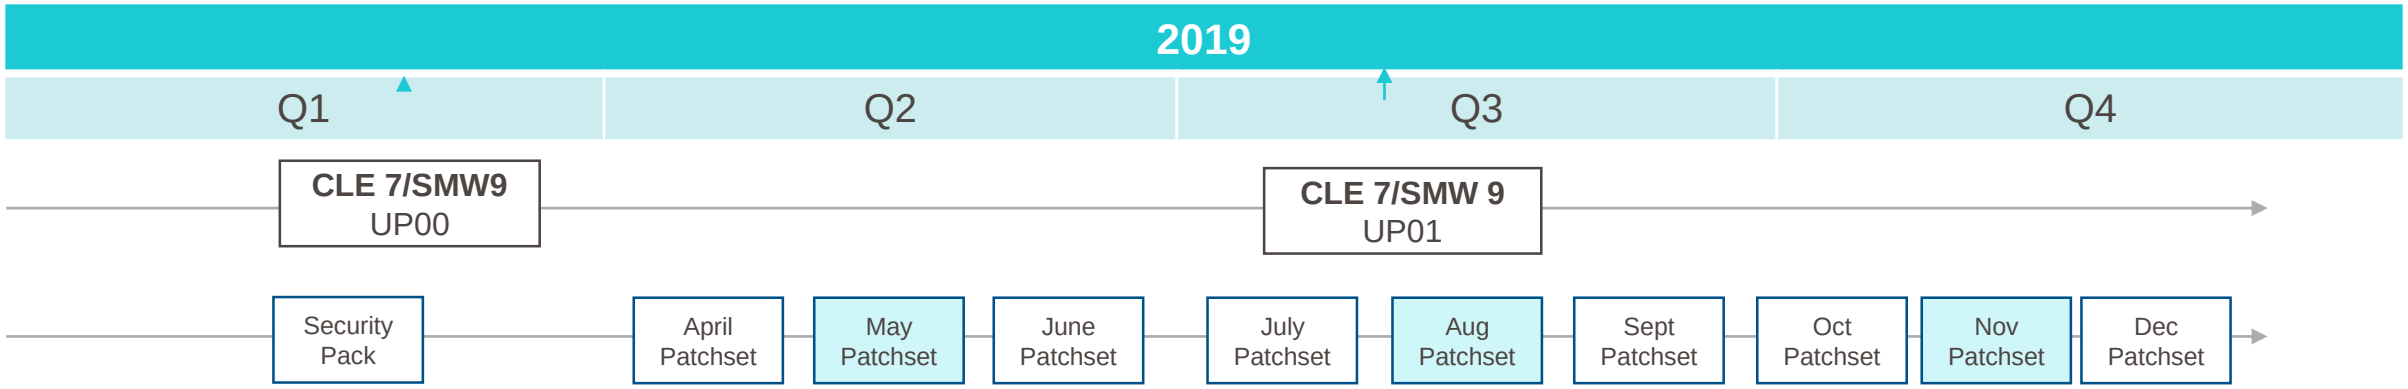

CLE 7.0/SMW 9.0 UP00

- SLES15
- Hugepage Memory reporting
- ARM B2 support
- CascadeLake
- NVIDIA Cuda 10.0
- Workload Managers w/SLES15 support
- Security Pack
- Lustre 2.11

CLE 7.0/SMW 9.0 UP01

- Bug Fixes
  - Logging
  - PE Image Push/sqpush
- Security Packs
  - 6.0.UP06
  - 6.0.UP07
  - 6.0.UP00
- Workload Manager qualifications
- NVIDIA Cuda 10.1
- Lustre 2.12

Key:

Patchset includes quarterly security pack

2020 XC ROADMAP – SOFTWARE

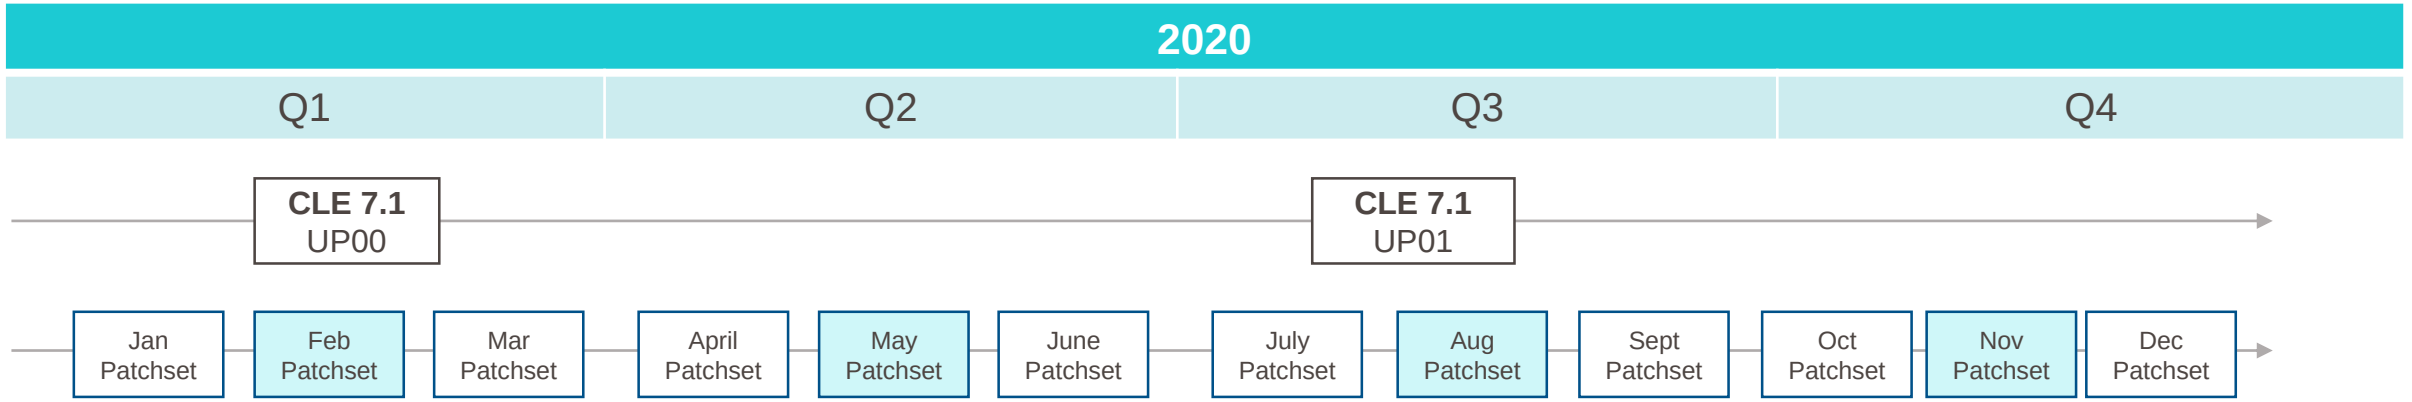

CLE 7.1/SMW 9.1 UP00

- SLES15 SP1
- Security Packs

CLE 7.1/SMW 9.1 UP01

- Bug Fixes
- Security Packs
  - 6.0UP07
  - 7.0UP00
  - 7.0UP01

Key:

Patchset includes quarterly security pack

XC Software Support Matrix

CLE/SMW Support Matrix

CLE/SMW Major or Minor Releases	CLE/SMW Update Releases	Release Date	*Full Support Ends	*Extended Support Ends	*Legacy Support Ends
CLE 7.0/ SMW 9.0 Final Major Release in support of XC systems	UP03	3Q30	TBD	TBD	N/A
	UP02	1Q20	CLE 7.0.UP03 GA	TBD	N/A
	UP01	3Q19	CLE 7.0.UP02 GA	8/27/2020	N/A
	UP00	2/28/2019	CLE 7.0.UP01 GA	2/27/2020	N/A
CLE 6.0/ SMW 8.0	UP07	7/12/2018	2/28/2019	2/28/2020	N/A All XC Platforms can upgrade to CLE 7.0
	UP06	3/1/2018	7/12/2018	8/28/2019	N/A
	UP05	10/5/2017	3/1/2018	7/12/2018	N/A
	UP04	6/27/2017	10/5/2017	3/1/2018	N/A
CLE 5.2/ SMW 7.2 Final Major Release in support of XE/XK systems	UP04 Final Update Release in support of XE/XK systems	09/25/2015	12/31/2017	1/31/2019	Ends 1/31/2022

Summary

- CLE 7.0 – performance and stability improvements
- Cumulative Patches – improved visibility for customers, improved quality
- Predictable roadmap cadence - 2 releases per year with monthly patches
- Look forward to more of your achievements on XC

SAFE HARBOR STATEMENT

This presentation may contain forward-looking statements that are based on our current expectations. Forward looking statements may include statements about our financial guidance and expected operating results, our opportunities and future potential, our product development and new product introduction plans, our ability to expand and penetrate our addressable markets and other statements that are not historical facts.

These statements are only predictions and actual results may materially vary from those projected. Please refer to Cray's documents filed with the SEC from time to time concerning factors that could affect the Company and these forward-looking statements.

THANK YOU

QUESTIONS?

kmark@cray.com

https://www.linkedin.com/in/kelly-mark

cray.com

@cray_inc

linkedin.com/company/cray-inc-/