



The Supercomputer Company


# XT3 - Scaling to New Heights: Status and Plans

David Wallace  
Cray Inc.

CUG 2006



This Presentation May Contain Some Preliminary Information, Subject To Change

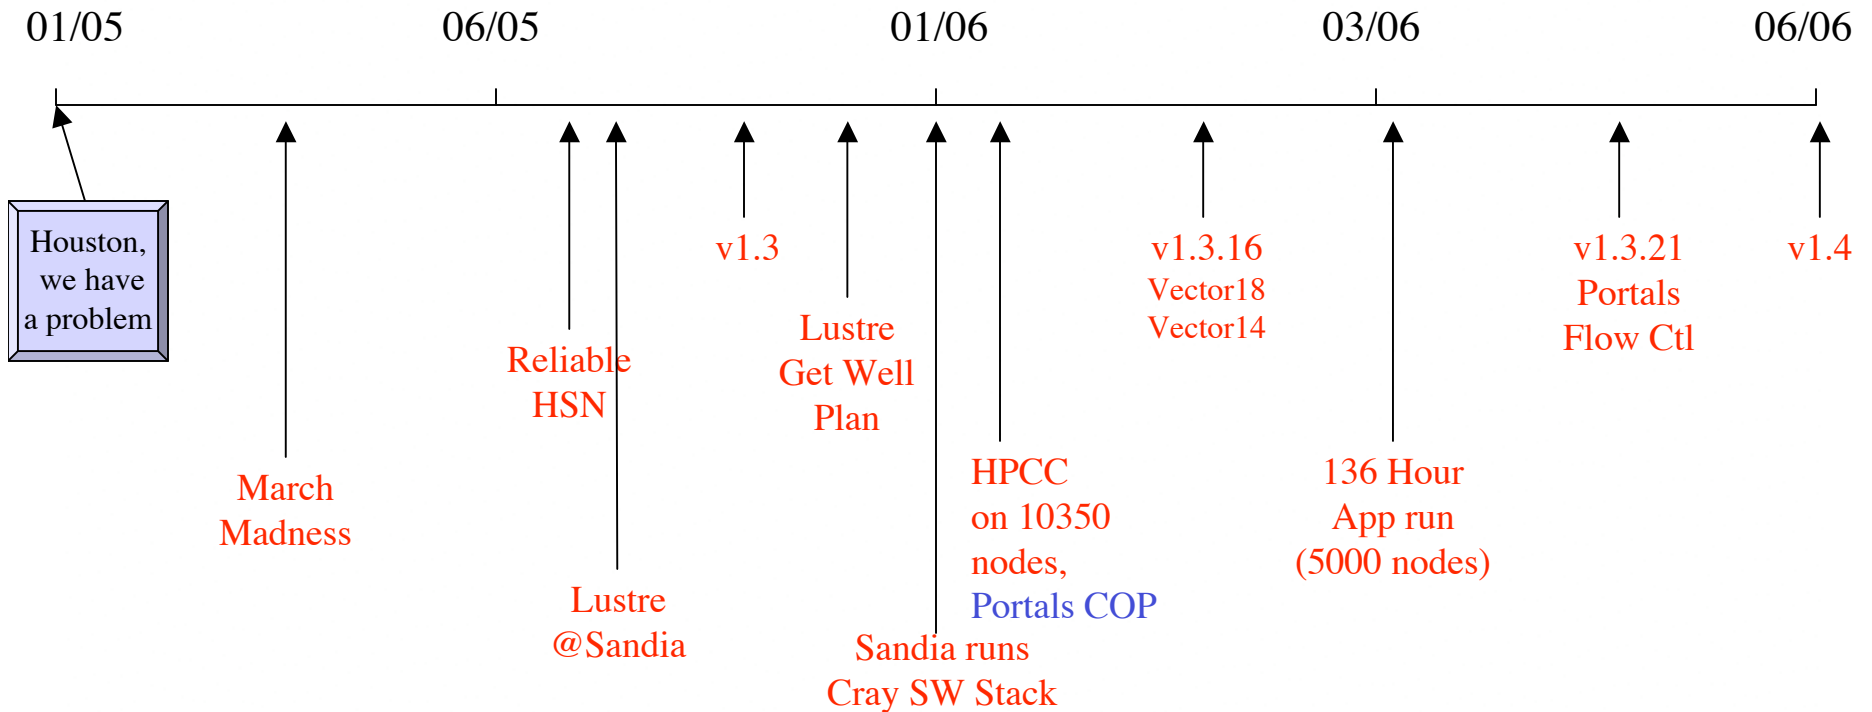


# Talking Points

- XT3: Major Milestones
- XT3: Software Support
- XT3: Looking Ahead

# XT3: Major Milestones

Many systems, many users, many applications



# XT3 has come a long way...

- Portals
  - Reliable scaling from about 0 nodes to 10000+ nodes
  - Latency reduction from 20+ to < 7us
  - COP and Flow control added for better network reliability and stability
  - Compute node warm boot
- Resiliency WRT node failures have improved
  - From almost nothing to “pretty good”
  - A year ago, pretty much any failure in the mesh would prevent boot and/or bring down the system
- Lustre “get well plan” very successful
- Many applications running
- Workload at customer sites building!


## XT3 Software Support

We still have a ways to go...

CUG 2006



This Presentation May Contain Some Preliminary Information, Subject To Change







# UNICOS/Ic: 1.3 Release Updates

- 1.3.07: Base release
- 1.3.09: Lustre, Portals, xtdumpsys
- 1.3.10: Portals, Lustre, Illegal Instructions
- 1.3.11: SMW/Diags, Kernel
- 1.3.12: Lustre performance, SMW/Diags, Features
- 1.3.13: Portals, Lustre, Security
- 1.3.15: SMW/Diags, OS, Qk
- 1.3.16: Qk, SMW
- 1.3.17: Portals, Lustre, Qk (**Regression**)
- 1.3.18: Portals, Lustre, Qk, MPICH
- 1.3.20: Portals, Qk, Lustre, SMW
- 1.3.21: Portals (flow control), CPA

# XT3 Closed SPRs (as of 4/30/06)




# XT3 Open SPRs (as of 4/30/06)





# XT3 SPR Comparison (as of 4/30/06)



# XT3 SPR Management

- SPR backlog reduction Plans:
  - Plan for Critical SPR interruptions
  - Schedule SPR work (non-Critical) into 2006 development plan
  - SPR week(s)
    - Stop feature development and focus on fixing SPRs
    - Various flavors
      - Focus on Critical and Urgent SPRs
      - Fix as many Major and Minor SPRs as possible
    - Need to plan for (re)stabilizing the DEV tree
    - Puts pressure on testing
  - Improve development process
    - Don't release defects!

**CRAY**

The Supercomputer Company


# XT3 Software Development

On-going Process Improvement

CUG 2006



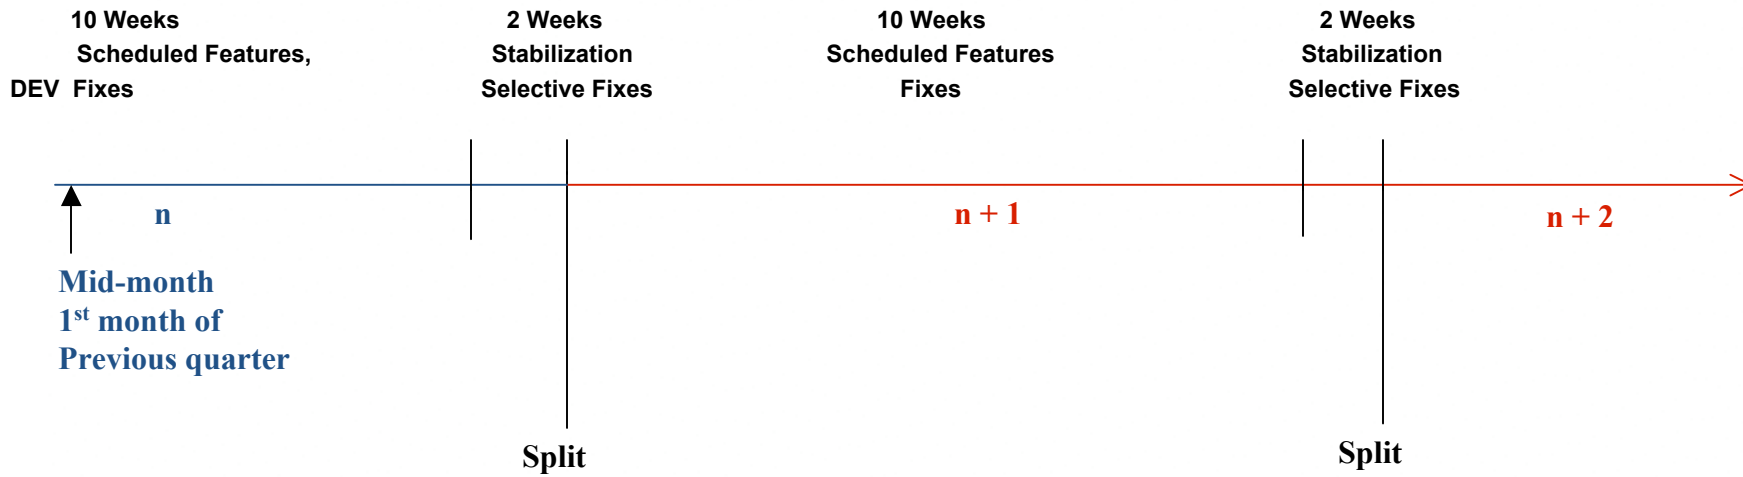
This Presentation May Contain Some Preliminary Information, Subject To Change



# Process Goals

- Quarterly releases allow for quarterly acceptances
  - Releases planned for the 2nd month of the quarter
  - Getting out quarterly releases, with only a six week release cycle, requires development to deliver working software at feature code cutoff
- Releases are supported, via updates (fix packages), with an overlap to allow time to upgrade
- 3 separate streams are maintained to allow for simultaneous development, release and support

# Development Stream



The 12-week cycle is broken into:

- Ten weeks of scheduled feature integration, and all levels of fixes; the most impactful features go in first
- Two weeks of stabilization, where only stabilizing fixes are taken

At the end of this 12-week cycle, code is split from dev to begin the release cycle; the next 12-week development cycle begins.

# Detailed Release Cycle

- Week 1: Build, Regression Test, Expose
- Week 2: Build, Regression Test, Expose
- Week 3: (Prerelease), Build, Regression Test, Expose
- Week 4: Build, Installation/Regression Test, Expose  
**Limited Availability Release**
- Week 5: Installation/Regression Test, Expose, Errata
- Week 6: Packaging, Final Errata, Media Duplication  
**General Availability Release**

Note: Sanity, Feature and some Regression Testing are done during the development cycle.



# Development Direction

- Ultimate goal is to have development source tree in a 'continuously releasable' state
  - Objective is ability to take snapshot of development tree, QA and release in less than 2 weeks
- Minimal destabilization of source tree
  - Changes controlled through Change Control Board review
  - Features and bugfixes subjected to peer code review and tested prior to integration into the development source tree
  - Daily build and testing to identify problems ASAP
- Partner with Customers to test on large systems (prior to LA release)

**CRAY**


The Supercomputer Company

# XT3: Looking Ahead

CUG 2006



This Presentation May Contain Some Preliminary Information, Subject To Change



# XT Roadmap

- Driving Forces –
  - XT Cleanup
    - Bugs – Reliability and Performance
      - Some bugs are just bugs and others require “refactoring”
  - New Hardware and Software commitments
    - Continuous stream of new hardware, devices, and customer requirements
  - New and Merged components for Common Systems
    - Have to move towards leveraged development

Real difficulty is keeping “forces” in motion and in concert

# XT3 Software Quality Improvement

- XT3 Software Quality continues to need improvement
- Steps taken last year have not resolved all problems
- Currently making additional changes to –
  - Improve the development processes for design and modifications
  - Improve bug reporting and fix reporting
  - Improving the checkout and review of changes
  - Changing the ownership and processes for integrating mods into release streams
  - Improving testing at all levels
  - Adding developers and testers to key areas

# XT3 Software Base Improvement

- Build and Integration of XT have not met evolving requirements
  - Working on a new build infrastructure –
    - Tools and processes shared with the rest of Software
    - Next steps are to break up the build into components
    - Separate unnecessary linkages
    - Provide a base for component releases...
- New requirements for Install are complete
  - Changes to build will provide a base for new install procedures
  - Install will be simplified and less prone to error
  - Install process will include capability of initial installs, upgrades, rollbacks, etc.



# XT3 Software Release Improvements

- Releases are currently being delivered quarterly with updates coming out on a weekly basis
  - Some delay in 1.4 because of scope of changes
  - Support for previous release will continue until GA of the next release
- Checkout of scaling for both features and fixes will become increasingly difficult as machine sizes approach 20K and 30K nodes
- We are discussing a model of release which would allow us to do testing, at customer sites, at scale, before the release
  - Need access to customer machines
  - Need changes to support rapid test and checkout
  - Need time to test and react to issues



# Scalar Platform High-Level Software Roadmap


Component	Rel 1.4 (2Q06)	Rel 1.x	Rel 2.x
HW Supported	XT3 Single and Dual Core	XT3 Single and Dual Core, XT4	XT3 and XT4
OS (SIO)	SLES 9	SLES 9	SLES 9
OS (Compute)	Catamount Virtual Node (CVN)	Catamount Virtual Node (CVN)	CVN and Compute Node Linux (CNL)
Runtime support	Yod/CPA	Yod/CPA	Yod/CPA or ALPS
HSN Protocol	Portals	Portals	Portals
File System	Lustre 1.4.6	Lustre 1.4.6?	Lustre 1.8
System Mgmt	CRMS	CRMS	CRMS
Compilers	PGI 6.1, gcc	PGI 6.x, gcc,	+PathScale
Tools	Apprentice <sup>2</sup> , CRAYPAT	Apprentice <sup>2</sup> , CRAYPAT	Apprentice <sup>2</sup> , CRAYPAT

## Questions

CUG 2006



This Presentation May Contain Some Preliminary Information, Subject To Change



**CRAY**


The Supercomputer Company

# Backup Slides

CUG 2006




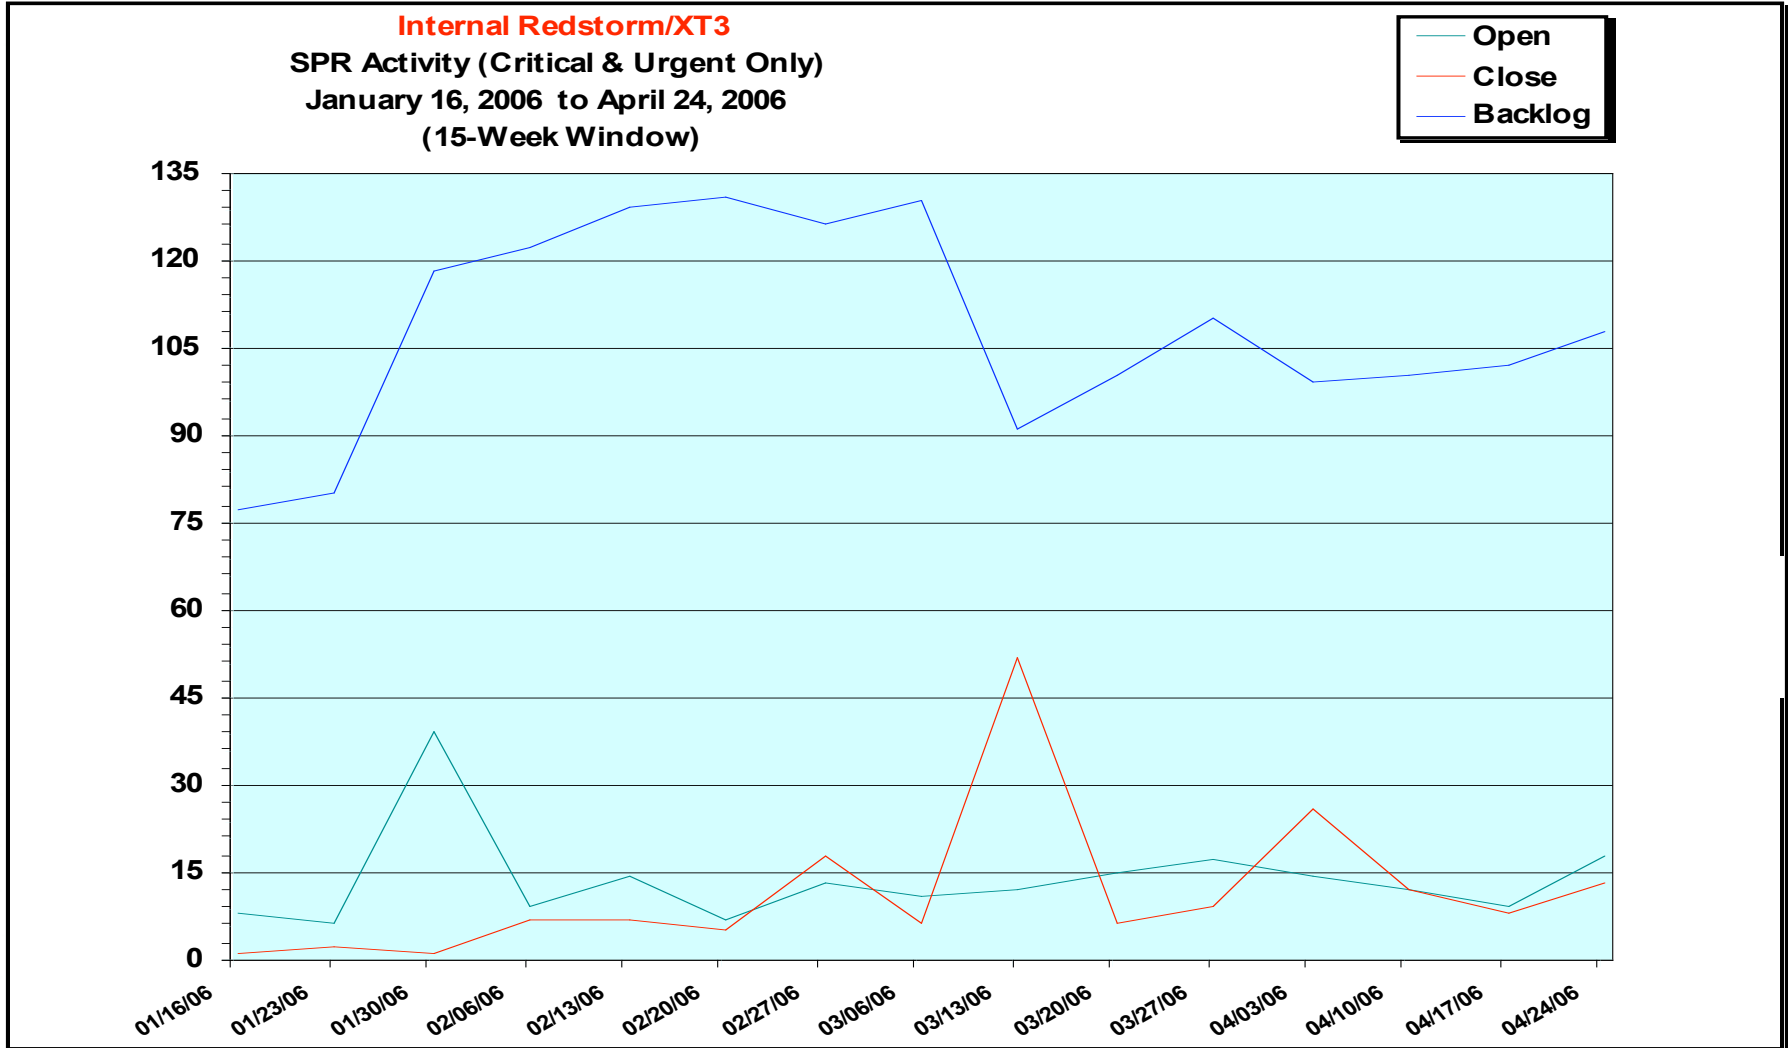
This Presentation May Contain Some Preliminary Information, Subject To Change



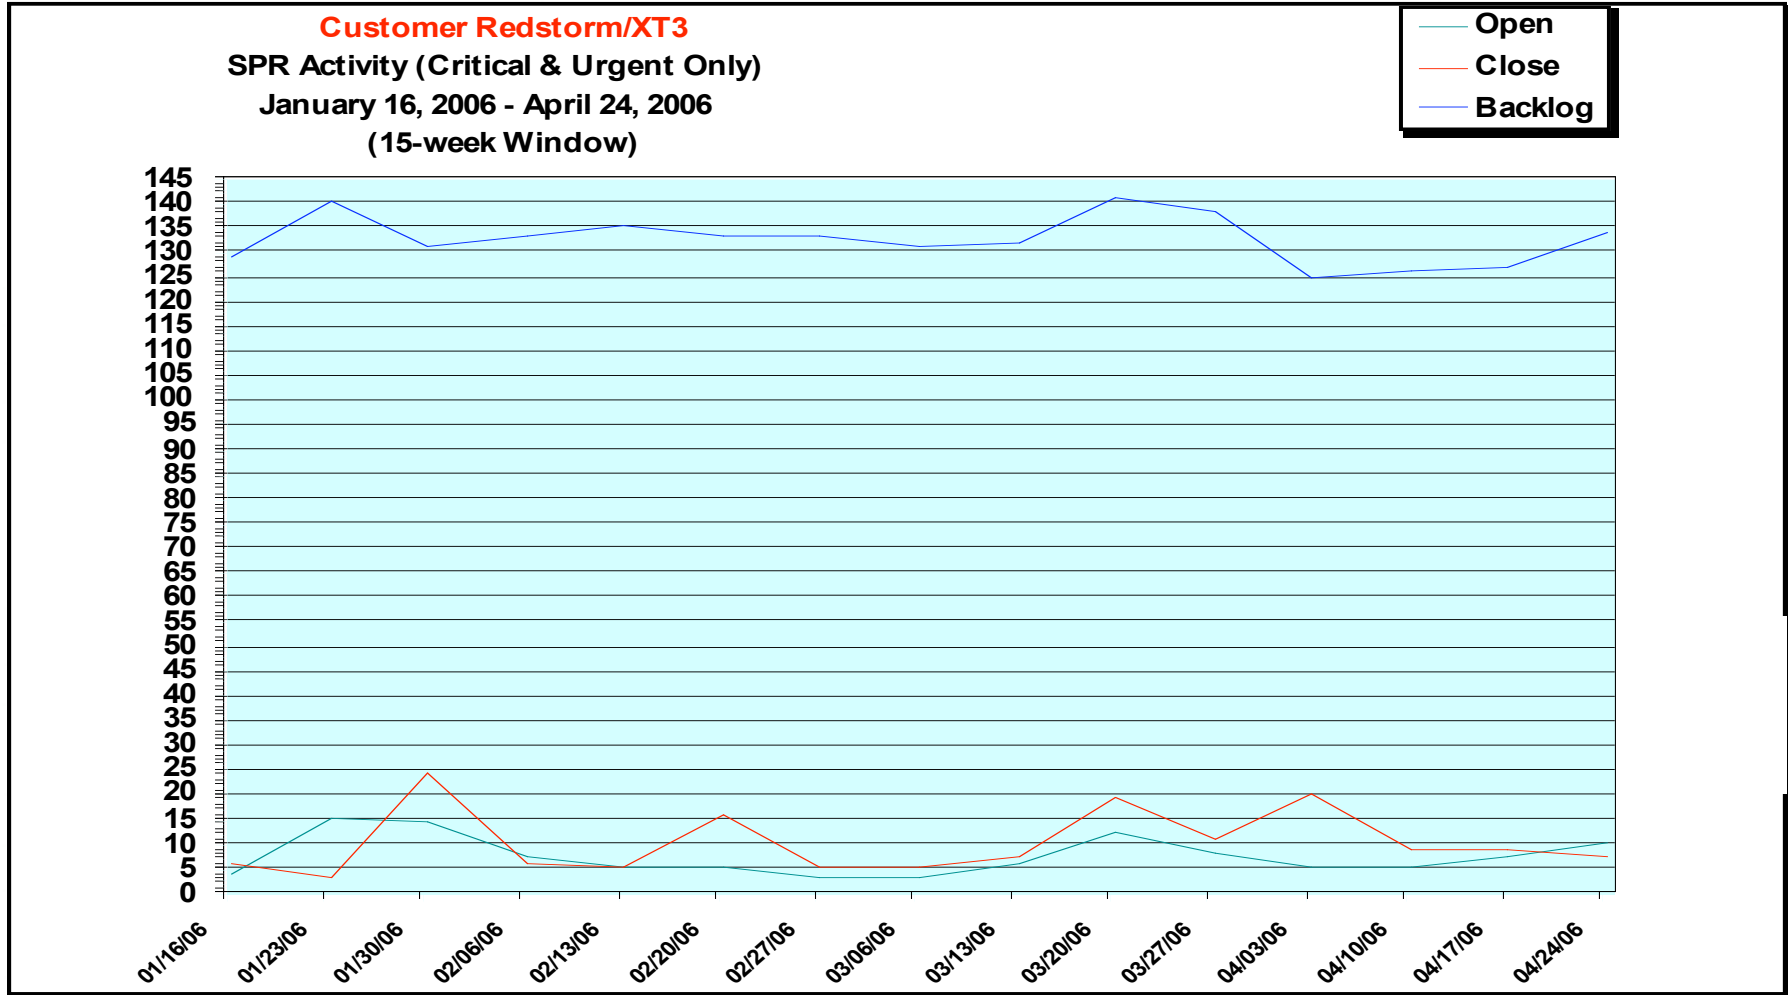
# Houston, we have a problem(s)...

- January 2005
  - Booting with the harness was unreliable at best.
  - Boot over HSN was not working yet.
  - RSMS was very young - barely working
  - Portals/FW "panic on error" (a feature) was in review. (old way: stop and freeze)
  - RX\_DMA\_EMPTY was our hot problem
    - light load crashes portals, and we didn't know why
  - Running "old" generic portals
    - Talking about "integrated" firmware, which would include accelerated.
  - Just starting to test MPI.
    - No apps to speak of were working.
  - The "harness" had been shipped to PSC, and they were trying to use it.



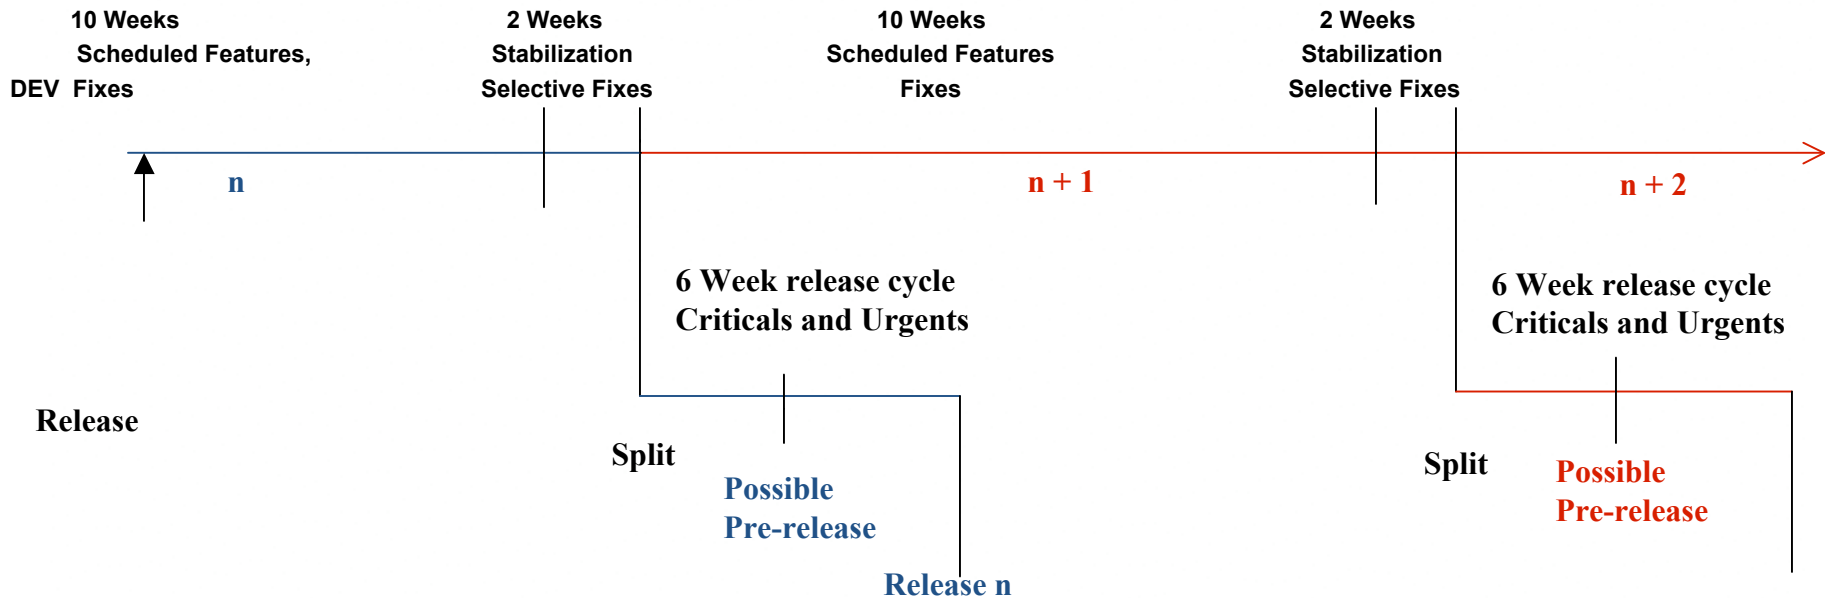








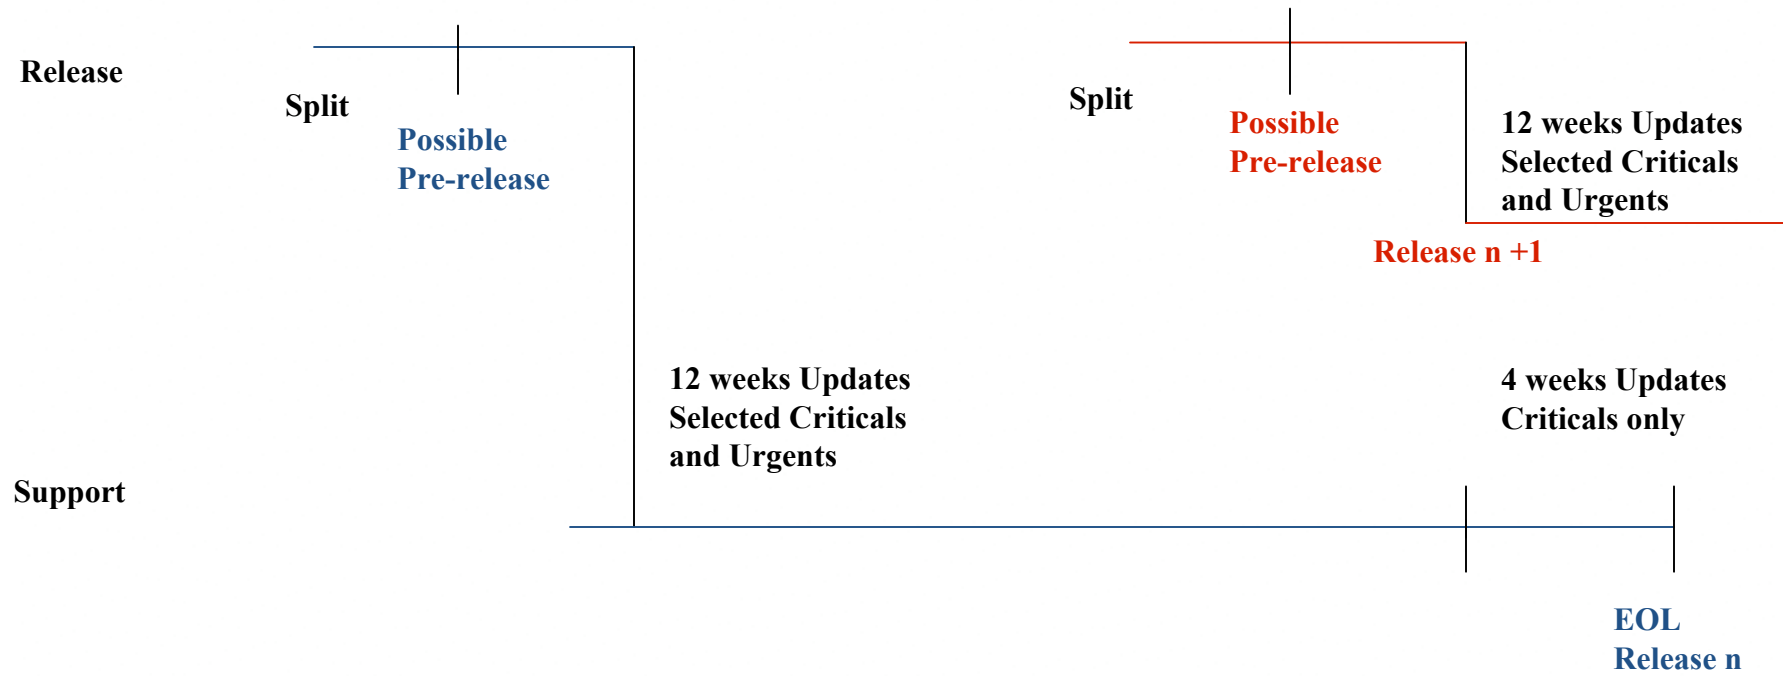
# Release Stream



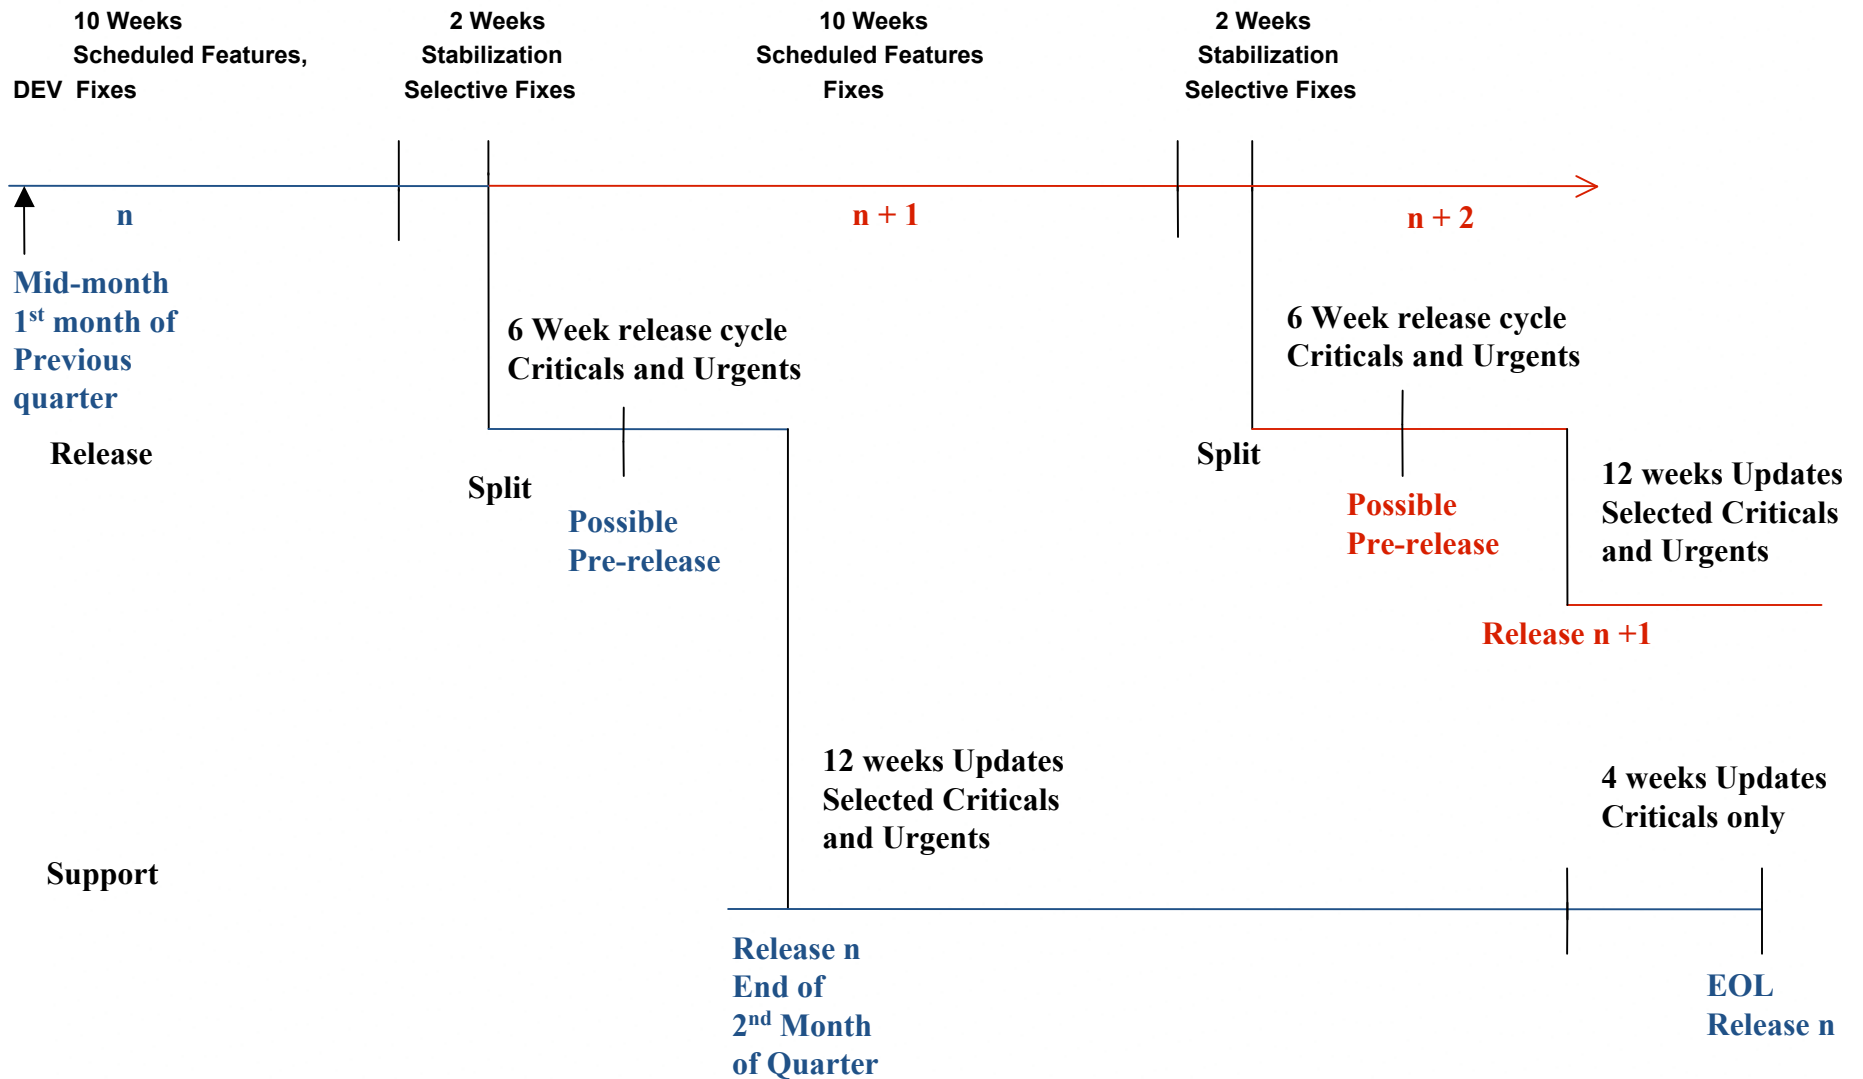
- After split, the release is tested and exposed during a six week release cycle
- Only fixes to critical and urgent problems are allowed
- The release occurs once release criteria have been met

# Support Stream

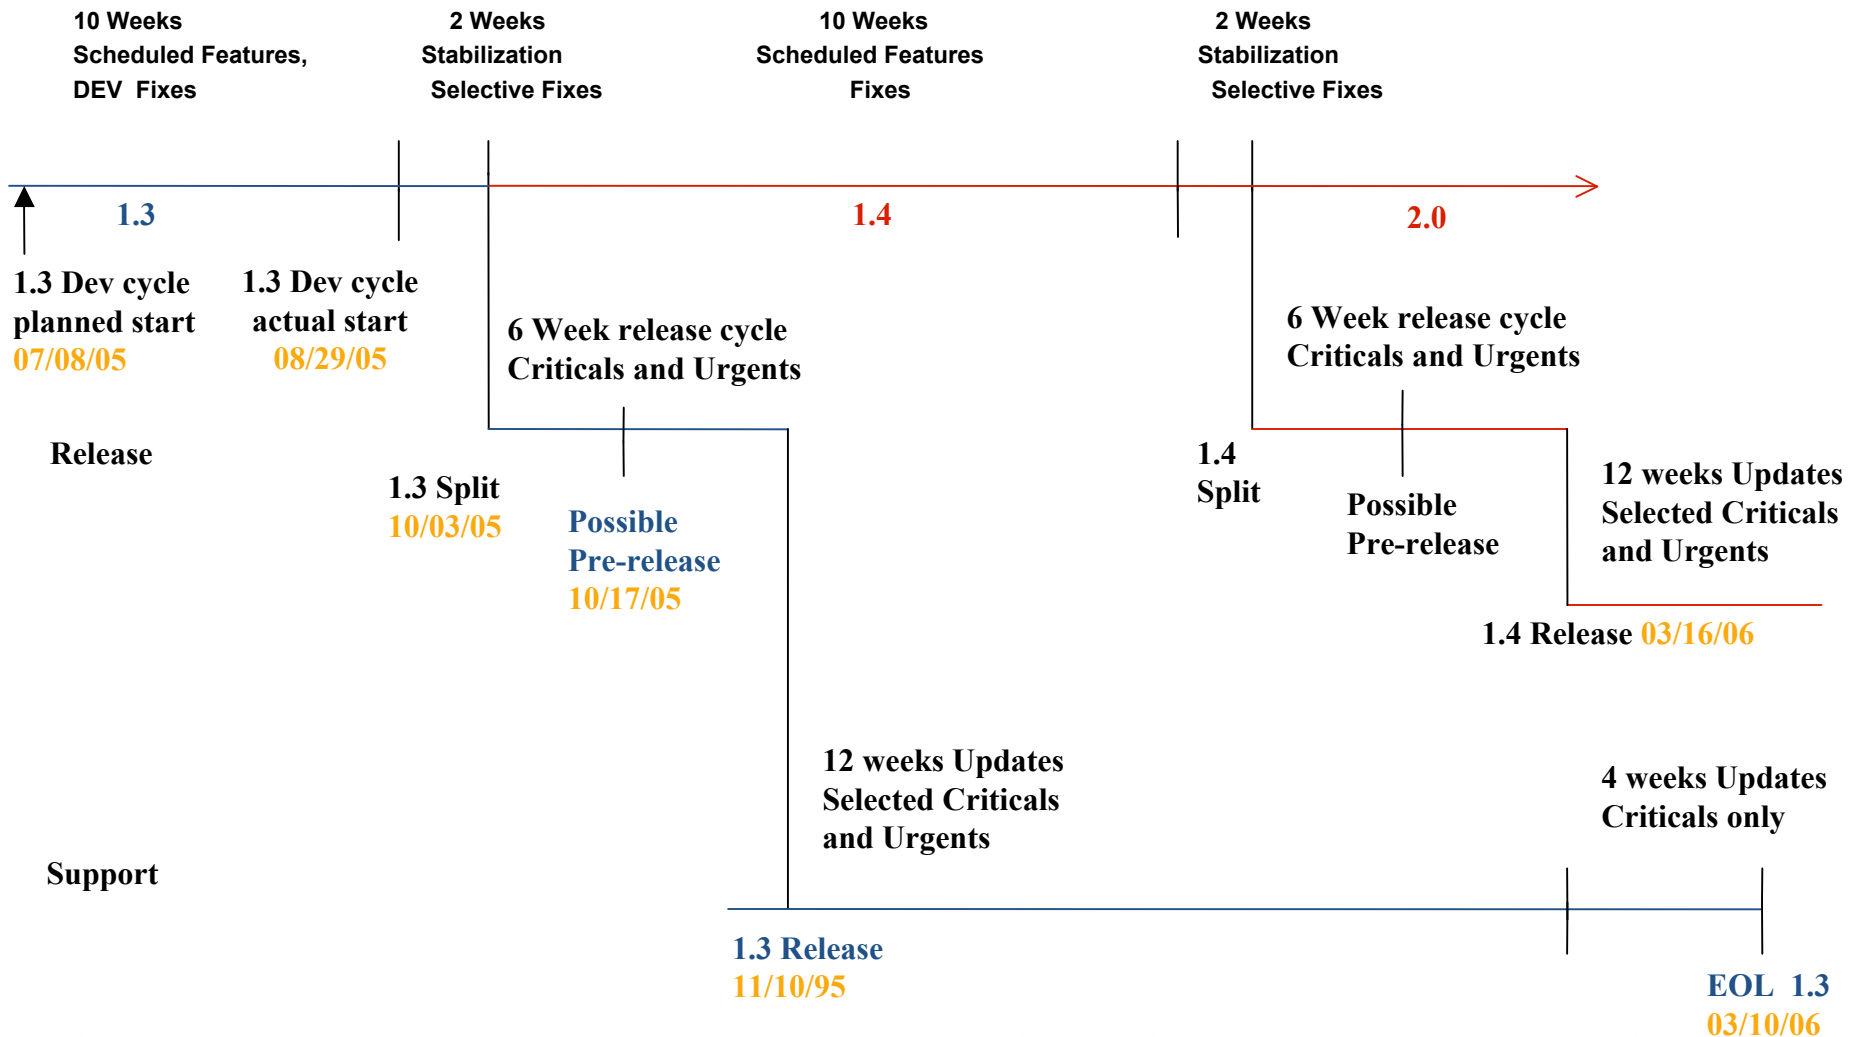
Updates (fix packages) occur up to weekly and contain fixes to selected critical and urgent problems for 12 weeks. Once the next release occurs, updates continue for 4 more weeks and include fixes to critical problems. This allows customers time to upgrade to the next release.



# Development, Release, Support Cycles



# XT3 1.3 Dev, Release, Support Dates



# Release Cycle Testing

