

# Sonexion Roadmap CUG 2015

Stan Friesen

Sonexion Product Manager

[sfriesen@cray.com](mailto:sfriesen@cray.com)



COMPUTE | STORE | ANALYZE

# 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.

---

COMPUTE | STORE | ANALYZE

# Sonexion Roadmap Content

- **Successes**
- **Challenges**
- **Lustre Strategy**
- **Sonexion Software Releases**
- **Product Platform Evolution**

---

COMPUTE | STORE | ANALYZE

# Sonexion Simple & Scalable System Architecture



- **Fully integrated rack**

- Prewired for InfiniBand, Ethernet and power
- Ready for instant expansion

- **Metadata Management Unit**

- Lustre metadata and management
- Systems monitoring and management

- **Scalable Storage Unit**

- Integrated object storage servers and targets
- GridRAID for fast rebuilds
- Balanced capacity and performance



COMPUTE | STORE | ANALYZE

# Cray Sonexion Success



- **Cray's flagship parallel file system product**
  - Sonexion 1300 in 2011 (3.0 GB/sec per SSU)
  - Sonexion 1600 in 2012 (5.0 GB/sec per SSU)
  - Sonexion 2000 in 2014 (7.5 GB/sec per SSU)
- **Over 1,500 SSUs & 90 ESUs (>175PB) in more than 125 file systems**
  - Over 40 customers
  - Globally accepted
    - UK Met installation underway in the UK
    - KAUST installation underway in Saudi Arabia
- **NCSA currently the largest file system**
  - 180 2TB Sonexion 1600 SSUs delivering:
    - 22 PB of usable capacity and over 1.0 TB/sec of sustained IO in 30 racks
- **Trinity will be installed in late 2015**
  - 216 6TB Sonexion 2000 SSUs delivering:
    - 80 PB of usable capacity and over 1.6 TB/sec of sustained IO in 36 racks

---

COMPUTE | STORE | ANALYZE

# Response to Sonexion Challenges



## ● Software Upgrades – too much down time

Challenges...

- Require significant and sometimes unaffordable amount of file system downtime
- Software upgrades are too complex – takes highly skilled expertise to perform
- Customers stay on older versions of software and don't get benefit of improvements

What we're doing about it...

- Near term goal is to reduce from days to hours (single rack systems <4hrs)
- Incremental upgrade improvements in 1.3.1 (infrastructure changes)

---

COMPUTE | STORE | ANALYZE

# Response to Sonexion Challenges



- **Support – takes too long to get fixes**

Challenges...

- Customer's backlog of bugs not getting into releases

What we're doing about it...

- Streamlining processes within Cray and with Seagate – staffing, single bug list
- Improve training to field support – 1<sup>st</sup> Line support training started
- Moved to single Lustre version for releases 1.3.1, 1.4, 1.5
- Reduce number of software versions in the field – converge on 1.5

---

COMPUTE | STORE | ANALYZE

# Response to Sonexion Challenges (Cont'd)



- **Installations – take too long**

Challenges...

- Requires additional upgrades and final preparation before being put online
- Performance often depends on site specific installation

What we're doing about it...

- Improved with software (1.4), documentation
- Improved Factory config. processes
- Goal is to ship directly from factory to site

---

COMPUTE | STORE | ANALYZE





- Seagate's Lustre distribution incorporates hundreds of stability and performance improvements
- Contributions from Intel, Seagate, and the community are curated to improve the quality of Lustre beyond the original OpenSFS release baseline
- **Lustre 2.1**
  - Converge on one Lustre version (1.3.1, 1.4, 1.5)
- **Lustre 2.5 (2.5.1)**
  - Seagate's Lustre 2.5 distribution is engineered to operate reliability at scale in production environments
- **Lustre 2.7 – future**
  - Cray and Seagate working with OpenSFS on improving future releases
  - Judicious in selecting, testing, packaging, and delivering the releases for Sonexion that are key to stable operation for Cray's customers.

# Sonexion Software 1.4



- **New features**
  - Lustre: 4MB IO, XATTR cache (Community contribution)
  - Dashboard, SSU/ESU Expansion, Guest Account
- **Lustre fixes/improvements**
  - Over 40 Sonexion bug fixes (quotas, flock deadlocks)
  - Over 330 community fixes
- **16 HA/Reliability Improvements**
- **20+ Infrastructure improvements (install/upgrade)**
- **~150 management and monitoring fixes/improvements**
- **Major USM update (2.16) with ~270 fixes/improvements**
  - IPMI/BMC, HW watchdogs, FW download
- **Available now**

---

COMPUTE | STORE | ANALYZE

# Sonexion Software 1.5



- **Major changes**

- Drive Drop Out fixes – a bundle of fixes that resolve the multiple drive drop out field issues
- Layered Stonith – Software to detect loss of physical access to the storage devices
- Advanced user authentication features (LDAP, NIS, passwd/group)

- **Latest USM – USM 3.2**

- **RAS (Reliability, Availability, Serviceability)**

- Guided FRU replacements for service personnel
- First ones: Disk Drive In Carrier (DDIC), Power Supply Unit (PSU)
- Roll out additional monitoring and as system updates

- **Availability**

- For Sonexion 2000 now
- For new Sonexion 1600s now
- For upgrades from 1.4 on Sonexion 1600s – June

---

COMPUTE | STORE | ANALYZE

# Sonexion Software 2.0



- **Lustre 2.5**
  - 868 patches (bug fixes) in Seagate Lustre
  - Plus new features
- **HSM Ready**
  - HSM integration with TAS
    - KAUST installation underway in Saudi Arabia
- **DNE Phase 1 support**
  - Hardware option. Support for ADU (Additional Directory Units)
  - UK Met Office installation underway in the UK
- **Availability**
  - For early customers on Sonexion 2000 now
  - Planning for support on Sonexion 1600
    - Date TBD based on Sonexion 2000 feedback & qualification

---

COMPUTE | STORE | ANALYZE

# Cray Sonexion<sup>®</sup> 2000



- **50% faster, 50% more capacity than Sonexion 1600**
  - 7.5 GB/sec/SSU, 2.1PB
  - Faster rebuilds by 3x
- **Improvements and new capabilities**
  - Data integrity on all drives
  - Windows access to file system
    - CNG – CIFS/NFS Gateway
  - Fast Ethernet connectivity (10/40GbE)
  - Metadata operations performance improvement (DNE with Lustre 2.5)
  - Cray TAS integration
- **Availability**
  - Product availability 4Q2014 with initial release of 1.5
    - Delivered a quarter ahead of schedule



COMPUTE | STORE | ANALYZE

# Sonexion 900



- **Infrastructure**
  - Modular architecture with similar building blocks as Sonexion 2000
- **All-in-One Metadata Management Unit (MMU)**
  - 2U24 drive enclosure with MDS and MGS
- **All-in-One Scalable Storage Unit (SSU)**
  - 4U24 drive enclosure with Object Storage Servers
- **Expansion Storage Unit (ESU)**
  - 4U24 drive enclosure with ESM modules
- **Does not require a Sonexion rack**
  - Can be housed in boot RAID rack, customer rack, etc.
  - No built-in InfiniBand switches
- **Scales up to 16 GB/s, 1PB usable capacity**
- **Available 1Q2015**



COMPUTE | STORE | ANALYZE

# Sonexion NextGen

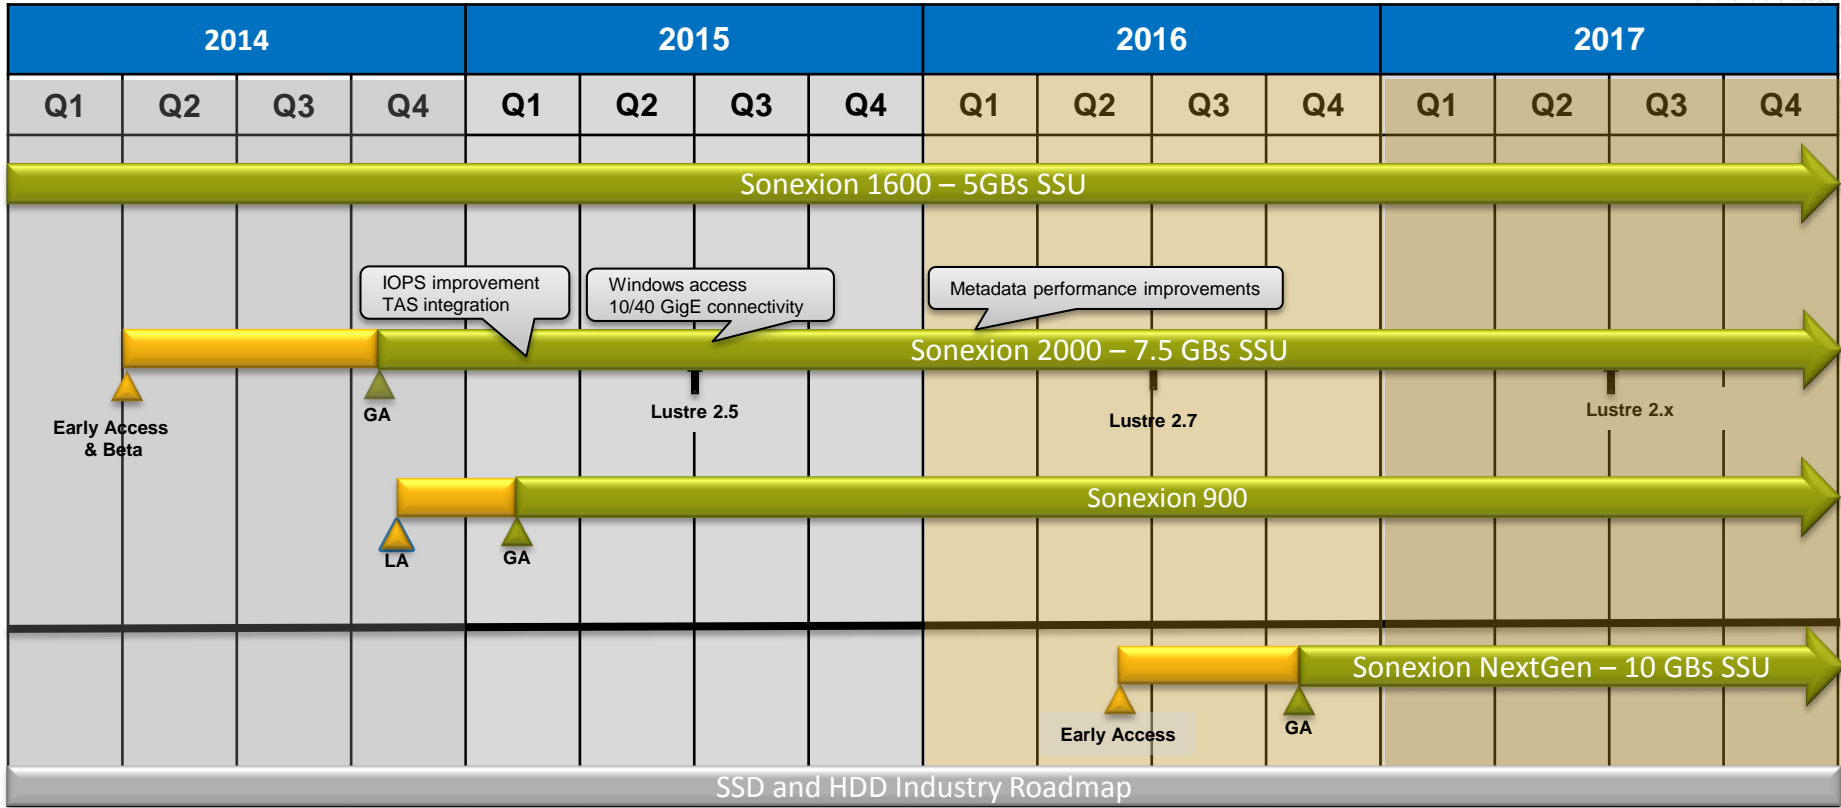


- Same form factor as Sonexion 2000
- 12Gb/s SAS
- Stronger servers
- Riding on a strong hard drive roadmap
- Availability
  - Product availability expected 4Q2016

---

COMPUTE | STORE | ANALYZE

# Sonexion Roadmap 2014 - 2017



COMPUTE | STORE |

Forecast only: uncommitted



# Q&A

---

COMPUTE | STORE | ANALYZE