ONStor™

Storage Solutions for the Content Generation

System-X Implementation Plan

A Story in Three Parts

Prologue

- The System X architecture is the key to taking this company from where it is now, a company that sells a few filers, to a serious threat to the biggest players in network storage.

- The implementation of SX, in order to be successful, must take into account many factors:

  a. The company cannot survive for two years without new products

  b. Available resources will be dynamic: resources will come and go due to support, maintenance, and new feature development needs

  c. A two year development project with no intermediate releases increases the project risk - plans and schedules are vulnerable and therefore less valuable to the company. Such projects have a low survival rate at startup companies.

Prologue, cont'd

Taking these considerations into account, the following general strategy will be used to prosecute SX development:

- A schedule of intermediate releases, some three or four, during a 24 month period
- Releases will be initially for existing hardware, possibly with some minor modifications
- Some new hardware designs (X86_64) will also be pursued during the 2 year time period
- Features and capabilities will be fed in as usual ala the traditional release train model

Part 1: Cougar2

"Use the Cores, Luke"

- Utilizing all 4 cores of 1480, able to serve CIFS, NFS, FTP, DCCP
- Replace TXRX runtime code with Linux and ...
- Implement new
  - Cluster Database (Replicating MySQL)
  - SCTP replacement of RMC
  - Port NFS and CIFS to Linux
  - Volume Manager
- Customer release in 12 months

ONStor™

Storage Solutions for the Content Generation

Part 2: NumaQugar

Single Linux kernel on both 1480 nodes

- Make Cougar a NUMA box (change memory addressing to be the same for both CPUs)
- Port our filesystem to Linux
  - A set of kernel threads?
  - Affinity grouped to the second 1480 node
- Multithread filesystem protocols
  - Group networking threads to first 1480 node
- Ditch 1125 processor and memory
  - Move CF controller to TXRX 1480
  - Add a dual GigE controller to TXRX PCI-X bus(?)

Part 3. HARDWARE

- Lose the separate Management processor: all-in-one architecture finally takes shape
- Do something Cougar-like with Opteron processor(s)
- Do something new-ish with 1480s
  - Three 1480s
  - Four 1480s?
- Eliminate CF and PCMCIA in favor of soldered down flash
  - Realistic life of product will not exceed life of modern flash chips
  - Access speed considerably faster
  - XIP possible
  - Built-in battery backup possible

ONStor™

Storage Solutions for the Content Generation

Resources

- Our resources for SX will be competing with other projects, especially those currently earmarked for Kegg.
- I believe Kegg needs to be scrapped, or re-defined from scratch.
- Much of original Kegg planning was done with the idea that it would be for Bobcat as well, which is no longer true.

Action Items:

- Pare down Kegg tasks to bare essentials - eliminating anything that will not carry over to SX or that is earmarked to be part of a subsystem that is to be replaced in Part 1. of SX
- Calculate resource earmarks
  - Exclusively assigned to SX: people that are to work on SX only
  - Primarily assigned to SX: people that are to work on SX unless responding to a interrupt from SW (System-W) or Escalations
  - Primarily SW: those resources not primarily assigned to escalations

Primary Koenigsegg Tasks

- Defect backlog: only defects that will be pertinent to SW
- Snapshot re-implementation
- Dump hardening - redesign and re-implement as necessary
- DM-IP and DM-SAN hardening
- HA failover flow audit/hardening (clustering flow)
- …

ONStor™

Storage Solutions for the Content Generation

Primary System-X Tasks

- New Clustering implementation
- SCTP based replacement of RMC (messaging subsystem)
- Integration of heartbeat and cluster manager
- New implementation of LUN discovery and Volume Management

Part 3: HARDWARE, cont.

- New uses for basic cougar design:
  - Front end caching box
  - Low power Web server appliance
  - Cluster switch/router (build in a switch with ports in back)
    - To be sold exclusively with Onstor chassis (not a general purpose network switch)
  - Load balancing front end (protocol level? session level? file level?)
  - Replication controller (8 GigE ports?)

ONStor™

Storage Solutions for the Content Generation