

PO LAN: Welcome; Wherefore?

Tim Chipman

July 31st 2006 Room 5603
Dalhousie University, Halifax, Nova Scotia

Overview

- Introduction / Why are we here ?
- Context: Brief background - TDC
- Goals
 - Establish clear communication, priorities
 - Assess & Document: Current Setup, “Needs” and “Wants” -> Develop plans / scenarios
- Work:
 - Current work underway
 - Possible work looming
 - Immediate issues to address today

Context – TDC Background

- “Local” origins:
 - Acadia 1994 Bsc Biology (Molbio)
 - Dal 1998 Msc Biochem (MolBio)
- Montreal Jaunt:
 - ~6mo. U.de.M MolBio Lab Research Associate
 - 1yr. McGill – General I.T. @ Comp.Centre
 - 5 yrs @ EcopiaBio: SysAdmin, then also I.T. Assoc.Director (team size range of 12->4..)

Context – IT @ Ecopia

- **Before:** Sun Solaris/Sparc based platform supporting core research operations. Major \$ investment in infrastructure, significant ongoing maintenance contracts
- **After:** Gradual introduction of Linux X86 platform to offload from Sun (computation cluster, thin client server, backup hardware, mail and web server, etc)
- **Benefits:** Lower capital costs, increased performance, lower operating costs

Goals for Today ... and beyond..

- **Stated Desired Endpoint:** Establish I.T. Infrastructure which supports research operations in the best possible manner, within the constraints of budget / staff available

(any additions / revisions?! :-)

- In order to reach this goal, very clear communication is required to properly understand requirements, solutions, and to develop a realistic plan (with short, mid, longer term scopes)

Current Platform (1/2)

- “Organic” network / server infrastructure evolution spanning ~15 years
- Core Server Infrastructure: Sun servers providing bulk of services (CPU intensive analysis, bulk data storage, backup, web-mail)
- Discrete Research Groups with resources partitioned (more-or-less) on per-group basis

Current Platform (2/2)

- “Not-Core” (not in server room) Server Infrastructure: some present in most research groups
 - evaluate: are there any ad-hoc “temporary” solutions present that might be improved?
- **Major Goal:** Evaluate current / looming requirements, develop plan for change to meet those needs, NOT simply to continue supporting all needs of past ~15 years.

Work Underway ^(1/3)

- Documentation of core systems configuration, minor optimization of configuration
- “Infrastructure Services Consolidation”:
(non-CPU-intensive -- but important roles!)
 - Mail Server
 - Public Website
 - ? Gateway server (inbound ssh, ftp, etc. from outside Dal Network)
 - ? Other options possible – using “virtual server” platform

Work Underway (2/3)

- Subsequently:

- Re-Deploy Current mail server as Sunray (thin client workstation) server
- Re-Deploy Current web server as Systems Monitoring Solution (?)

- Storage Related Topics:

- Complete Disk Use Audit
- Deploy unused array
- Cleanup of “Forgotten TempFiles”
- Consolidate / Rearrange “split storage pools” to “unified storage pools”
- Evaluate growth trends, develop plan for future

Work “Underway” (3/3)

- Cargo Deployment (Sun v20z - dual-2.4ghz opteron with 1.8 Tb disk array)
 - Never deployed, should be rolled out as part of “Anchor transition” plan(~4.8 vs Anchor's 6 ghz)
 - Target role:
 - General data analysis server role (as per Anchor)
 - Migrate users & all data from Anchor SAN-> Disk Array
 - Add: Unified Samba/NFS File Sharing for easier access
 - Possibly: deploy as Linux (CentOS) to facilitate Matlab (not available for Solaris-x86 platform)

Possible Looming Work ^(1/4)

- Water into wine: How to improve things without spending any \$?
 - Limited by:
 - Well defined priorities
 - Time
 - “Best” (free/cheap) solutions for highest priorities

Possible Looming Work (2/4)

- **Increased Network Security:**
 - Block all inbound traffic to POLan from outside Dal campus (with some exceptions):
 - Public web server
 - Gateway server for SSH/FTP
 - Grant access from Dal VPN subnet for all current traffic
- **Monitoring Server:** (sysadmin use only)
 - Server loads, utilization; network traffic
- **Backup Solution:** (L20 robot, open source software)

Possible Looming Work ^(3/4)

- **Intranet Web Content Structure**
 - internal communication structures ?
- **Improved Fileserver** for Group Collaboration
 - unix, windows access for "not large data" files (word docs, etc)
- **Intranet DMS** (Document Management Server)
 - improved collaborative work within/between groups
- **"Tech Support" Web-Email Ticket System**
 - for clear task assignment / resolution / tracking /
"FAQ" Database development
- **"Tech Support" Wiki** - "Reference /HowTo" docs
- ...etc....

Possible Looming Work (4/4)

- Server Room Fun:

- New Rack to Rackmount Core Gear (Jade, NAS Array, 2 Disk arrays) – (not free: costs \$500-\$1000)
- Retire under-utilized (or unstable) hardware following appropriate “service migration” as appropriate
 - SGI Origin? Anchor ? Other ?

...pending deployment of offload server(s)...

- ❖ NB: Anchor = 8 x 750mhz = “6000mhz”
- ❖ 2-way Dual-Core Opteron 1u Rackmount Server (4 CPU cores@ 2ghz) = “8000mhz”
- ❖ Cost (CASTALIA) ~ \$4-5k per box

Unknown / Inbound Work ..! ?

- **GOAL:** TDC to provide “I.T. consulting / planning” **advice** on appropriate research grant spending **in advance** to ensure good fit for:
 - “New solution requirements”,
 - “PO Lan context” and
 - “Long term plan” for overall environment
- **Examples:**
 - Inbound Beowulf / Computation Cluster...
 - Fall/2006-Target: NSERC Equipment Grant Submission
 - Other New Research Projects ?
 - Aside: Possible benefits of virtualized servers / consolidated hardware)

Issues to Address “right now”

- Discussion on above topics ...?
- Schedule 1-on-1 meetings ...?
- I.T. Policy Development before Sept...?
 - Proposal: UCIS material as starting point to establish well defined “line in sand” for acceptable use, etc. which SHOULD evolve to meet the requirements of the PO_Lan environment

All Done!

- Further topics for discussion ?