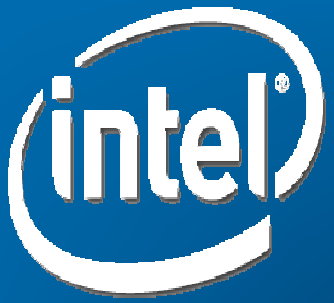

Improving User Productivity in a Cloud Environment

April 5, 2012

Kiron Pai

Design Automation Lead

Converged Core Development Organization, Intel Corp.

Cloud Computing

Scorecard

KPI	Last week			5 weeks average		
	Actual	Target	Status	Actual	Target	Status
Run Time Failure Rate	18.28%	17%		15.12%	17%	
Job Slot Utilization	99.84%	85%		99.82%	85%	

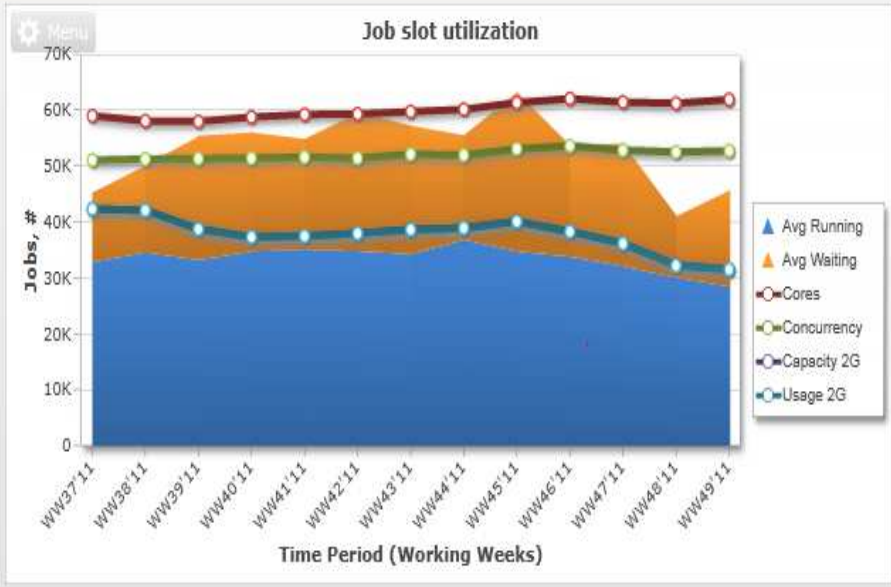

Computing optimized for Bulk Processing

Interactive Computing (ideal)

Improve runtime

Improve Reliability

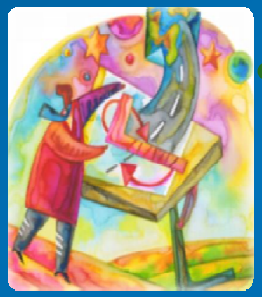

Enable Creativity

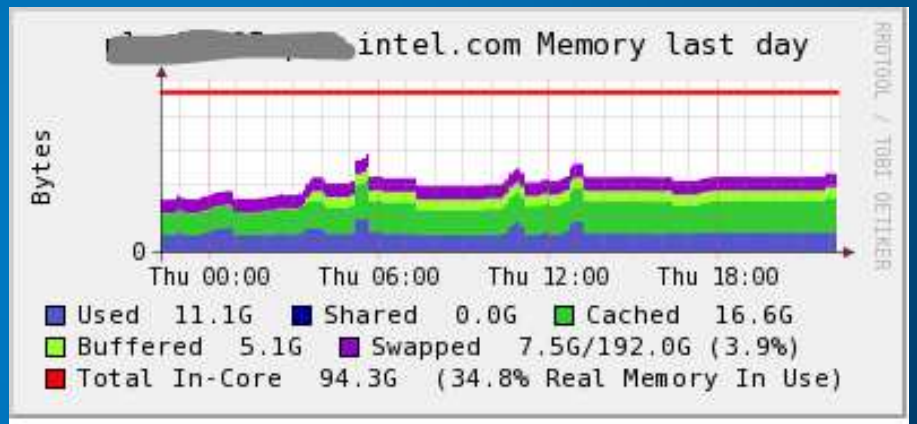

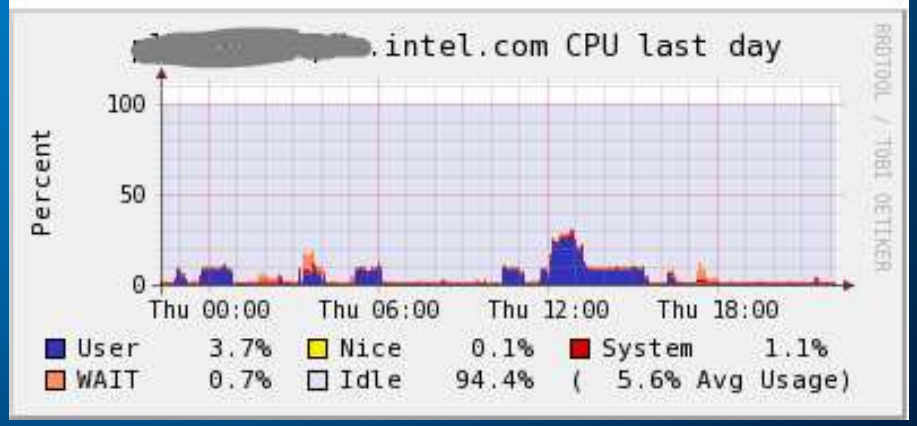

Computing optimized for User Experience

Interactive Computing (reality)

Unpredictable Runtime

Frequent Slowdowns

Computing adds to chaos

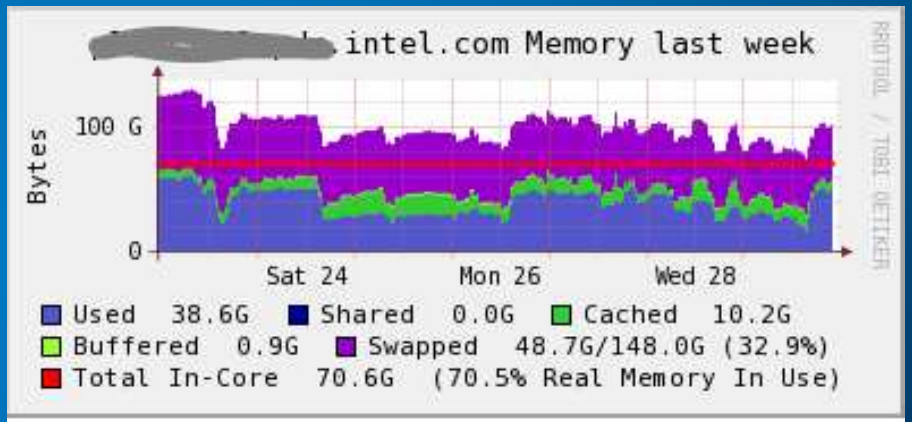

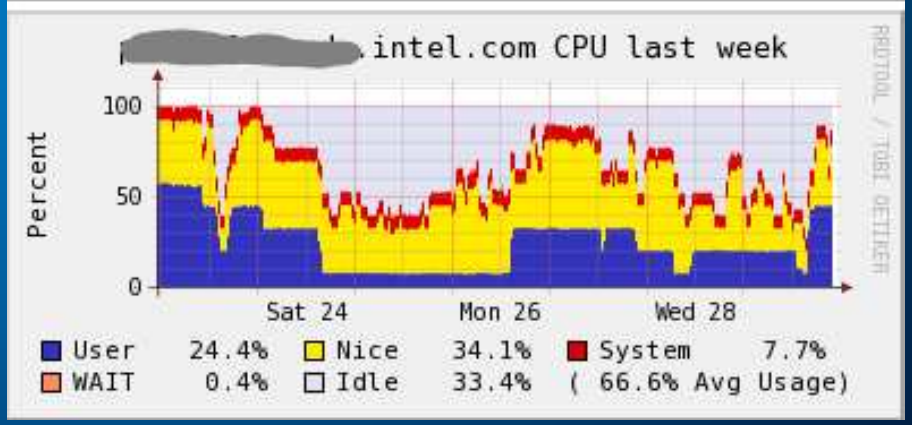

Compromised Experience

Goals

- Improve user productivity by engineering an interactive compute environment
  - Improve run time
  - Improve reliability
  - Neutral on total cost of ownership (TCO)

Interactive Computing Solution of Today

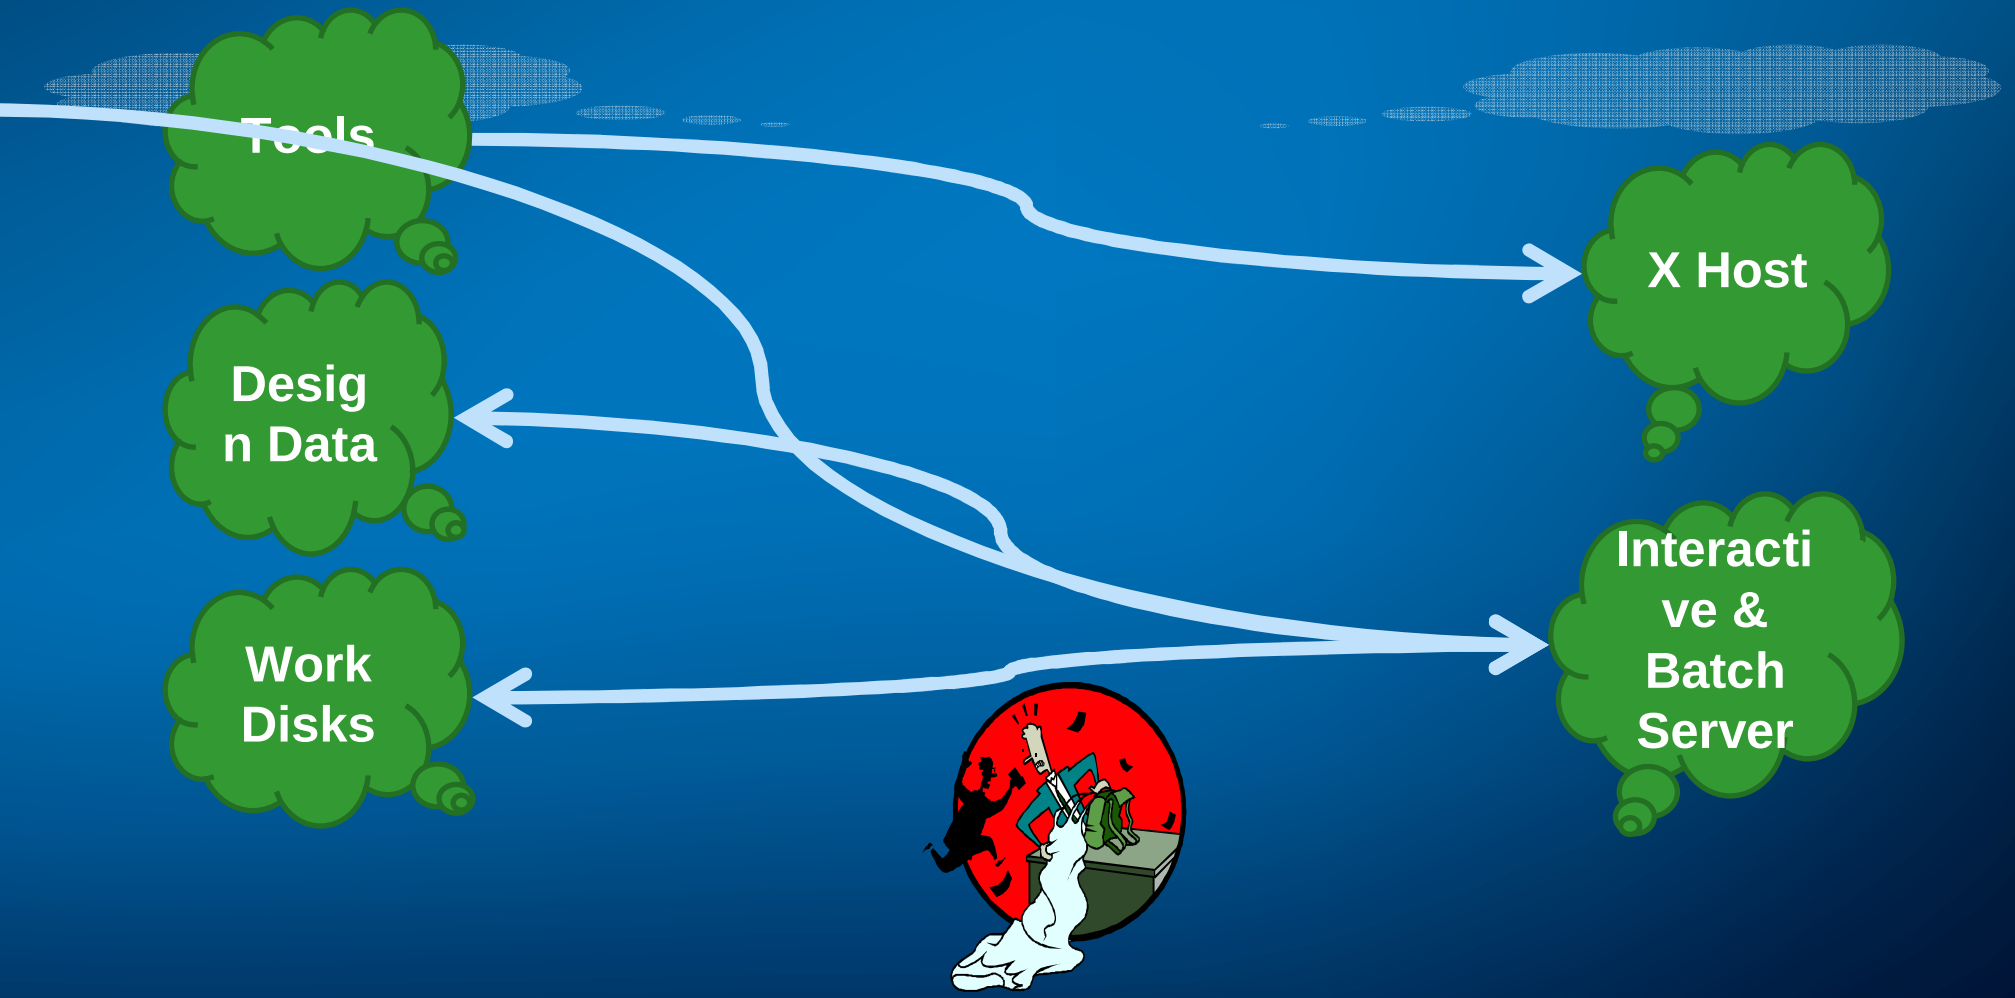

Storage

Compute

Interdependent Solution: Unpredictable

Solution

- Improve Runtime
  - Allocate one high end machine per user
  - Use local disk for work disk and tools
- Improve Reliability
  - Reduce multiple points of failure
    - Shared VNC servers
    - Multiple users/batch jobs per compute servers
    - Multiple fileserver
- Reduce total cost of ownership
  - Unused cycles used by batch
  - Low Cost fileservers for work disk backup

Interactive Computing Solution of Future

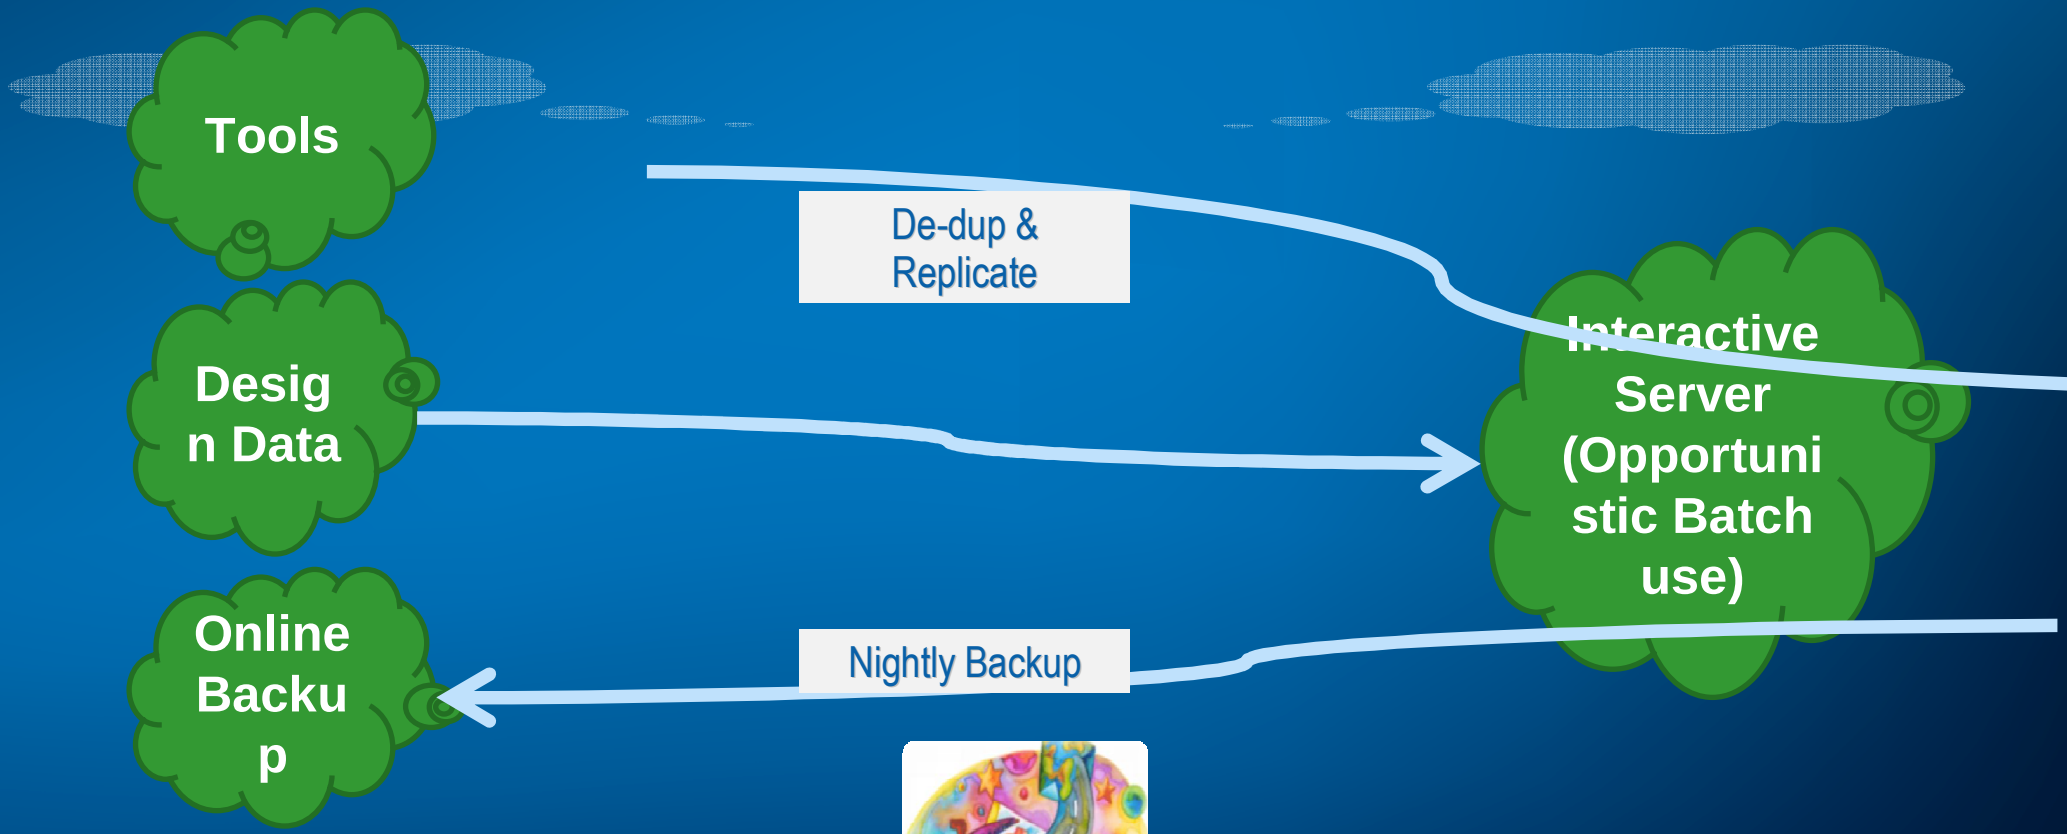

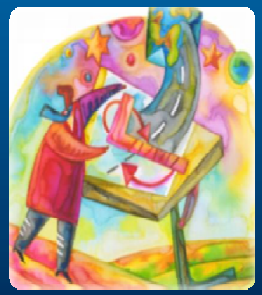

Storage

Compute

Isolation Solution : Runtime, Reliability Improvement

Results/Progress Summary

Solution	Goals	
	Runtime & Reliability Improvement	TCO
One user per m/c	+ Reduced thrash due to multiple users + Headroom on m/c for changing usage model	- Increases cost of compute per user
Local disk work area	+ Runtime improvement ~ 20% + Reduces dependency on Tier 2 work disk storage	
VNC on local machine	+ Reduced hops for VNC + Reduced interruptions due to shared machine slowdowns	
Opportunistic batch for unused cycles		- Mixed results, works only for low profile use
Low cost storage for online backup		+ Reduced cost of work disk for users
Mirror tools to local disk	- de-dup on tools disks shrank size to 600G only (orig. est. was <300G) + Caching method designed, reduces space need ~100G - High risk – complex solution, changes core pieces of design environment	

Results/Progress - Runtime

[redacted] -- All runs were on dedicated system

	NHM 2.93GHz, 48GB (C) Data: NFS Workarea: NFS	NHM 2.93GHz, 48GB (B) Data: NFS Workarea: SAS, ordered	NHM 2.93GHz, 48GB (C) Data: SAS, ordered Workarea: SAS, ordered	NHM 2.93GHz, 48GB (D) Data: SAS, writeback Workarea: SAS, writeback
Stage2 - Synthesis (DC)	4:19:29	4:01:08	3:59:27	4:00:39
Stage3 - ICC Place	3:11:56	2:40:01	2:32:25	2:33:00
Stage4 - ICC clock	5:05:51	4:09:49	3:53:06	3:49:17
Stage5 - ICC Route	3:26:05	2:48:37	2:44:18	2:43:01
Stage6 - Sizing	11:40:55	11:31:21	11:11:34	10:54:53
Stage6 - Optimization	11:54:05	11:22:49	11:23:37	10:47:32
Stage6 - DFM/Backend	3:18:47	2:38:11	2:22:10	2:23:06
Stage7 - Verification	23:43:49	22:17:28	22:00:55	21:22:11
Total Runtime	66:40:57	61:29:24	60:07:32	58:33:39
Scalability	BASE	1.08	1.11	1.14
		BASE	1.02	1.05

[redacted] -- All runs were on dedicated system

	NHM 2.93GHz, 48GB (C) Data: NFS Workarea: NFS	NHM 2.93GHz, 48GB (B) Data: NFS Workarea: SAS, ordered	WSM 3.06GHz, 48GB (C) Data: SAS, ordered Workarea: SAS, ordered	WSM 3.06GHz, 96GB (D) Data: SAS, writeback Workarea: SAS, writeback	WSM 3.06GHz, 96GB (D) Data: SSD, CFQ/writeback Workarea: SSD, CFQ/WB
Stage2 - Synthesis (DC)	1:08:25	1:08:14	1:01:52	1:00:33	1:00:39
Stage3 - ICC Place	1:33:42	1:15:18	1:07:55	1:06:50	1:07:45
Stage4 - ICC dock	1:38:38	1:17:46	1:09:25	1:06:37	1:08:14
Stage5 - ICC Route	1:43:52	1:24:49	1:20:35	1:18:37	1:19:26
Stage6 - Sizing	5:09:49	4:24:28	4:10:29	4:09:50	4:07:04
Stage6 - Optimization	4:01:14	3:43:51	3:29:26	3:27:58	3:27:13
Stage6 - DFM/Backend	2:14:55	1:45:10	1:33:49	1:35:28	1:31:16
Stage7 - Verification	8:34:38	7:47:33	7:44:13	7:39:57	7:32:32
Total Runtime	26:05:13	22:47:09	21:37:44	21:25:50	21:14:09
Scalability	BASE	1.14	1.21	1.22	1.23
		BASE	1.05	1.06	1.07

SOURCE: Ty Tang, Ananth Sankaranarayanan, Kripa Sankaranarayanan (Intel Corp)

Clear runtime benefits of 1 user/mc & use of local disk

Results/Progress - Reliability

- Detailed tracking of CPU, memory, localdisk i/o, network
  - Separate out user and batch
- Track events when thresholds crossed
  - Frequency and duration
- User perception

Measurable improvement

Results/Progress - TCO

		Cost per user per year
Current (shared) model – 6 users per machine (12C/96G)	Compute	$ 250
	Avg. space per user of 250G Tier2	$ 375
	Avg. batch utilization of 10%	- $ 25
	Total	$ 600
Proposed Solution – 1 user per machine (12C/96G)	Compute	$1500
	Avg. space per user of 250GB Tier4	$ 180
	Avg. batch utilization of 60%	- $ 900
	T2 Tools fileserver (Tools caching)	- $ 75
	Total	$ 705

TCO goal is not met, opportunistic batch is key to reducing cost

Next Steps

- Improve monitoring methods
- Improve batch job policy
- Improve tools caching

Backup

Why de-dup?

- Frequently released data has large number of duplicated content
- De-duplication resolves duplicates to unique file using hard links
- Improves
  - Reduces space usage (less frequent cleanup)
  - Improved sync (faster/reliable)
  - Reduced support

de-dup - Some Stats

Tool	Size before dedup (GB)	Size after dedup (GB)
rls	19	0.5
fm	15	0.11
process	11	1.25
extraction	9	0.21
explorer	8	1.3
clktools	6	1.11
netlister	5	1.64
	85	6.12

Why replicate tools to local disk?

- Significant runtime benefit
- Reduced reliance on network and tools fileservers
- peer 2 peer replication for reduced reliance on central fileservers

References

Data DeDuplication:

http://en.wikipedia.org/wiki/Data_deduplication

PigeonHole Principle:

http://en.wikipedia.org/wiki/Pigeonhole_principle

Achieving Scalability and Availability with Peer-to-Peer Transactional Replication:

http://msdn.microsoft.com/en-us/library/cc966404.aspx