ITAK

Information
Technology
Asset
Knowledge

Volume 9 Issue 11

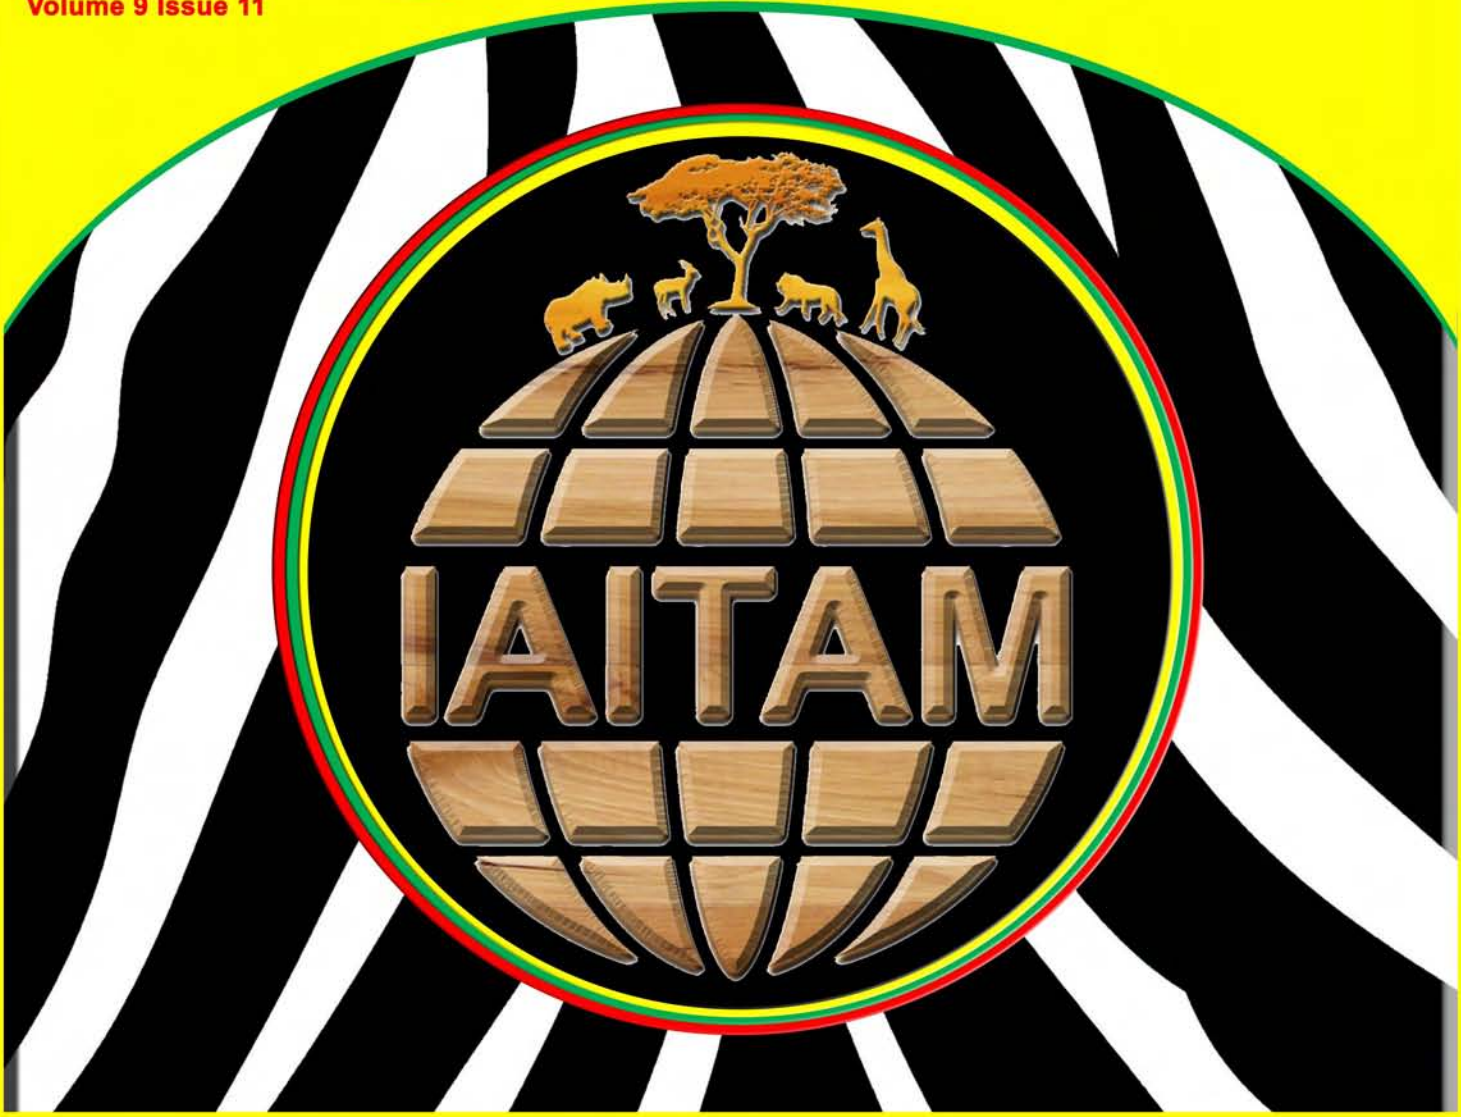

IAITAM 2014
FALL CONFERENCE

Benefits of Usage Metering in Cloud Solutions

Measuring and Managing in the Cloud

By: **Linda M. Cole**
Director for Sales, North America, ***Open iT, Inc.***

Introduction

Cloud has become a huge buzzword, from the C-suite to an organization's end users, because of the significant benefits cloud can provide to big and small businesses across a host of industry sectors.

According to a Gartner study, cloud usage is growing exponentially and over half of large enterprises will have cloud deployments by 2017. In fact, by 2017, IDC expects public IT cloud services to drive 17% of IT product spending and nearly half of all growth across five technology categories: applications, system infrastructure software, platform as a service (PaaS), servers, and basic storage. Software as a service (SaaS) is expected to remain the largest public IT cloud service category throughout the forecast, capturing 59.7% of revenues in 2017.

The shift to the cloud has essentially provided a great way for companies to collaborate and access files anywhere and anytime. This is especially true for the managed cloud environments that allow the implementation of tools that run automated management and reporting systems, helping organizations to meter and maximize the use of IT resources such as software licenses.

With the cloud, software vendors are changing their license strategies and pricing. Many companies these days are paying their cloud provider for server and seat licenses. Since software is such a huge cost, vendors are continuously challenged to prove their product's value and justify their pricing.

Furthermore, managing compliance has never been easy. The potentially complex task of accurately assessing the number of licenses used has plagued many companies due to inconsistent licensing terms being used nowadays. Many software vendors do not have a tool to measure the actual usage of their software. Therefore, audits are still the main tool for monitoring compliance, a burden for customers.

Predictability is valuable in business, especially when assessing the aggregate corporate requirements for software applications that are required to accomplish corporate goals and objectives. Optimizing IT assets effectively depends on an accurate understanding of historical application use. The best way to achieve optimized assets is through usage metering tools. Software usage tools measure how much and how often software assets are used.

Why Meter Cloud?

There are many facets to cloud computing and the term cloud means different things. Unfortunately, just like ITAM, there is no "one tool for everything." Today, there are tools that measure different aspects of cloud computing. Some provide daily emails regarding spending data, alerts to sudden changes, spikes and trends. There are tools that measure various aspects of overall system performance.

The value of metering is in the ability to correlate the context of the usage to the business need. Many tools report on a single metric. However, a true metering tool allows users to analyze all of the data in different ways to give management a true story of what is actually happening and provide them the ability to make more informed decisions.

Cloud providers bill by usage and there are many stories about "runaway" cloud bills that could have been avoided or minimized with the use of appropriate metering and reporting tools. Here are just a few reasons companies may not achieve maximum ROI on their cloud spending:

1. **Over-provisioning:** IT typically builds to on-premise standards instead of current use. The benefit of cloud is to add resources as needed. On-premise requires capacity planning for 3-5 years.
2. **Forgetting temporary resources:** Test Dev Servers, short term projects, etc. may be provisioned and forgotten.
3. **User error in provisioning:** Lack of knowledge or experience in provisioning may have the IT administrator selecting the wrong package.
4. **Not understanding license implications:** IT administrators must know how applications are licensed. The IT administrator may increase capacity in the cloud platform from the on-premise apps and cause larger costs in licenses such as Oracle or IBM, which licenses on capacity. Or, the administrator may put up an SQL instance on a public facing server that now requires a different type of license with a much different cost structure.

Benchmarking resources before cloud deployment allows better upfront provisioning to minimize excess capacity. It is important to remember that resources are being billed even if they are not being used.

Cloud Billing Challenges

Many cloud providers for IaaS bill for resources provisioned but not used, for example CPU, RAM, I/O, or storage. IT administrators provision the individual components the same way they would the hardware specs for a new physical or virtual server on-premise in their environment. Traditionally, when listing hardware specs for a server on-premise, one considers what the server will be used for and plans for a minimum of a three year life span, which means that one plans for excess capacity. In the cloud environment, the specs are for today's requirements and simply add resources as needed.

Cloud providers advertise elasticity and the ability to adjust resources up and down. However, many providers sell packages that require 1-, 2- or 3-year commitments and they are happy to have users add additional resources, but they do not allow for reductions. So, companies with seasonal or trending needs must look carefully at the parameters of the cloud provider's business agreement.

Another difference is in the thought process and historical definitions, especially around PPU. In a software PPU agreement, inactivity means no charge during inactivity. However, in a cloud IaaS, inactivity does not typically stop the billing. There are tools to spin down servers when not in use where the configuration is not lost, but IT administrators must have a process to ensure the consistent use of this approach. This is especially important in dev test environments, where developers want to only spin up images while the actual testing is being performed but can inadvertently leave these running in an idle state.

Metering Reports for Server Resources

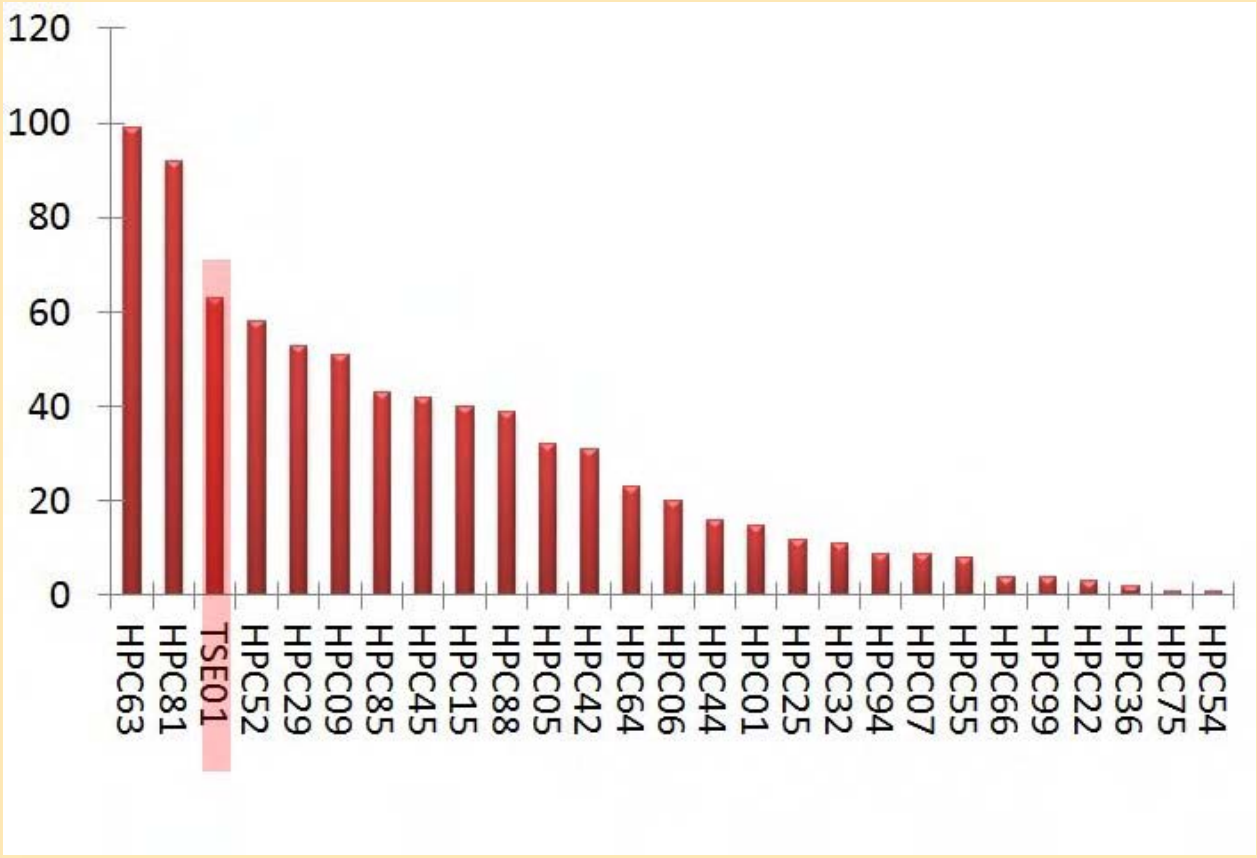

Figure 1: CPU % Utilization

Figure 1 shows the CPU utilization as a percentage, that is, the relationship between the used and the total CPU available in seconds per month. The load is very unevenly distributed among the host computers. Many computers have less than 10% load, while others have more than 90%. Most of the nodes are HPC nodes and have high resource utilization. However, TSE is a terminal server meant for interactive jobs, where the load is typically kept at less than 10%.

Some might look at this report and just assume more resources need to be added. However, before deciding what to do, it is not only important to know how much CPU utilization the various nodes have, but also to know or analyze what the computers are doing. Why does the machine HPC63 have such a high load?

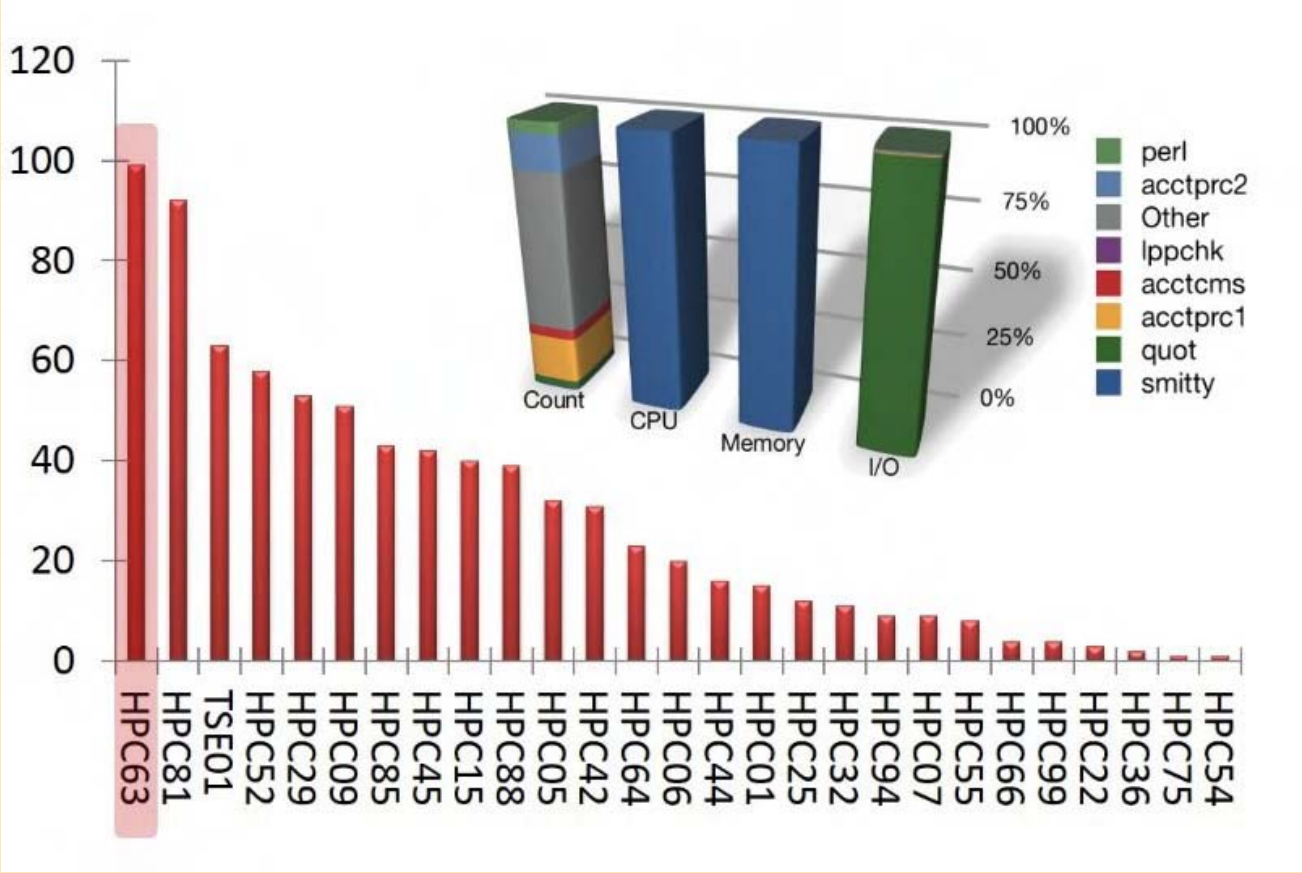

Figure 2: CPU - Usage in Percent – HPC63

Figure 2 shows the details of what has taken place on HPC63 during the period. Here, we can see right away what has happened. SMITTY, a system admin tool on AIX, has had a problem. It is apparently stuck in a loop and is consuming substantially all of the CPU. This makes the host unavailable for any other process to use. Therefore, it is important to detect this type of problem as soon as possible. This should have been discovered quickly, but in large organizations with many machines, it may be easy to lose sight of each machine, especially if they do not have tools in place with error condition detection and alerting mechanisms.

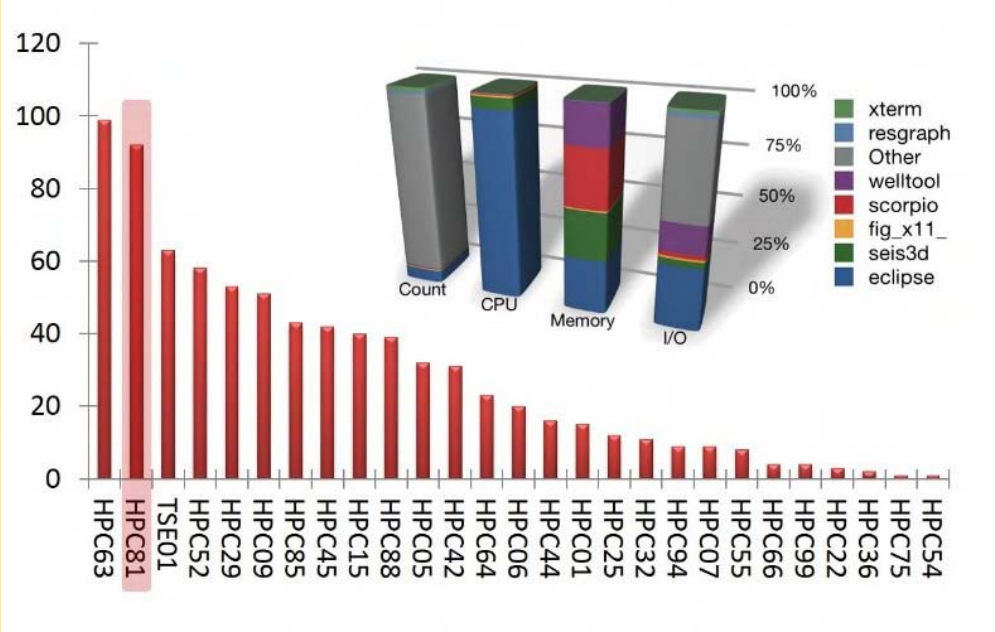

Figure 3: CPU - Usage in Percent - HPC81

In Figure 3, by utilizing metering tools and drilling down into the host HPC81, it can be seen immediately that the eclipse app in this example has consumed 93% of the CPU and 25% of the memory.

Looking at the table data that generated this report, it can also be seen that the program has been started 536 times last month. Therefore, it would be simple to reconfigure the cluster management so that eclipse would load more evenly across multiple nodes.

Metering Reports for Storage Resources

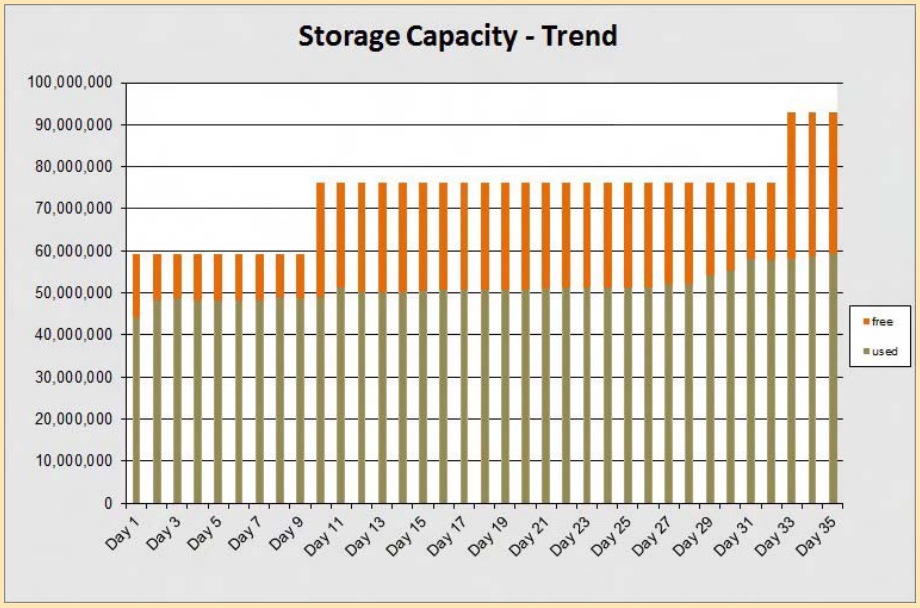

Figure 4: Storage Capacity Trend

In Figure 4, the report shows the total amount of disk storage capacity and how it divides between free and used storage space. When free storage space declines, IT administrators buy more disks. However, it is important not just to look at the bulk storage figures, but also to look at what is consuming the disk space.

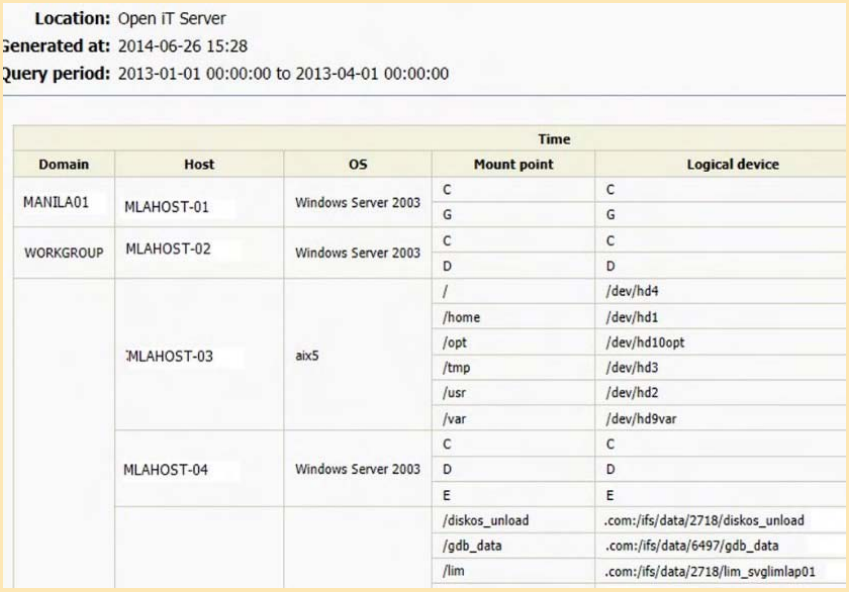

Location: Open iT Server
Generated at: 2014-06-26 15:28
Query period: 2013-01-01 00:00:00 to 2013-04-01 00:00:00

			Time	
Domain	**Host**	**OS**	**Mount point**	**Logical device**
MANILA01	MLAHOST-01	Windows Server 2003	C	C
			G	G
WORKGROUP	MLAHOST-02	Windows Server 2003	C	C
			D	D
	MLAHOST-03	aix5	/	/dev/hd4
			/home	/dev/hd1
			/opt	/dev/hd10opt
			/tmp	/dev/hd3
			/usr	/dev/hd2
			/var	/dev/hd9var
	MLAHOST-04	Windows Server 2003	C	C
			D	D
			E	E
			/diskos_unload	.com:/ifs/data/2718/diskos_unload
			/gdb_data	.com:/ifs/data/6497/gdb_data
			/lim	.com:/ifs/data/2718/lim_svglimlap01

Figure 5: Disk Space by Host and Share

The report in Figure 5 breaks down the used space by host and share. Most IT groups can still easily generate reports like this, showing usage by share folder. However, by combining that information with the temperature of the files being stored, more information can be seen and deduced as shown in Figure 6.

Server
-27 00:19
-01 00:00:00 to 2013-04-01 00:00:00

Host	File system	Account	File type	File status	Read temperature	Modify temperature	Number of files	Logical size	Physical size	Reported cost
mnl001win	C	<empty>	CAD	plain	tepid	frozen	18.00	2.46	2.46	284.28
			animation	plain	tepid	frozen	6.00	0.47	0.47	54.74
			archive	plain	frozen	frozen	6.00	0.02	0.02	1.93
					tepid	frozen	6.00	1.56	1.56	179.77
			audio	plain	hot	frozen	10.00	2.51	2.51	289.21
					tepid	frozen	588.00	49.24	49.24	5,679.82
					warm	frozen	2.00	0.04	0.04	4.47
			backup	plain	frozen	frozen	6.00	0.00	0.00	0.07
					hot	frozen	6.00	0.12	0.12	14.37
						hot	10.00	0.01	0.01	1.51
					tepid	cold	6.00	1.29	1.29	148.95
						frozen	7.71	0.00	0.00	0.12
						tepid	6.00	0.00	0.00	0.16
					warm	warm	2.00	0.01	0.01	0.76
				hidden	tepid	frozen	12.00	8.25	8.25	951.58
				plain	cold	cold	93.00	0.00	0.00	0.08
					frozen	frozen	123.00	1.77	1.77	203.96
					hot	cold	90.00	16.18	16.18	1,866.12
						frozen	1,117.00	550.57	550.57	63,504.15
						hot	8.00	0.00	0.00	0.00
						tepid	464.00	117.71	117.71	13,577.13
						warm	60.00	56.30	56.30	6,493.89
						cold	585.00	170.87	170.87	19,708.61

Figure 6: Storage by Type/Temp/Size

The report in Figure 6 shows file type, read temperature, modify temperature, number of files, etc., which makes this an actionable report. So, the "frozen" (defined as no access in 3 years) can be found and moved to Tier 3 or other less expensive storage. Similar reports that show file type by age or file owner by age may also aid in disk cleanup. Also, there may be benefits in backup as well since static files do not need to be backed up daily.

Showback vs Chargeback

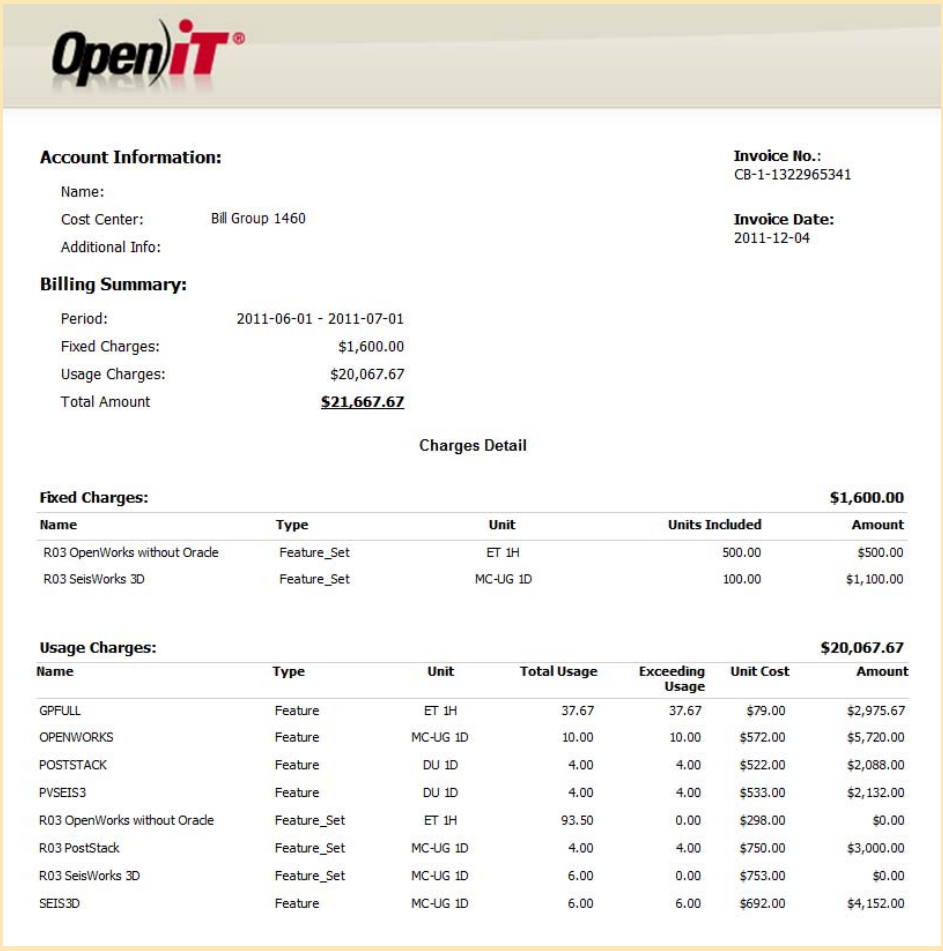

OpenIT

Account Information:

Name:

Cost Center: Bill Group 1460

Additional Info:

Invoice No.: CB-1-1322965341

Invoice Date: 2011-12-04

Billing Summary:

Period:	2011-06-01 - 2011-07-01
Fixed Charges:	$1,600.00
Usage Charges:	$20,067.67
Total Amount	**$21,667.67**

Charges Detail

Fixed Charges: **$1,600.00**

Name	Type	Unit	Units Included	Amount
R03 OpenWorks without Oracle	Feature_Set	ET 1H	500.00	$500.00
R03 SeisWorks 3D	Feature_Set	MC-UG 1D	100.00	$1,100.00

Usage Charges: **$20,067.67**

Name	Type	Unit	Total Usage	Exceeding Usage	Unit Cost	Amount
GPFULL	Feature	ET 1H	37.67	37.67	$79.00	$2,975.67
OPENWORKS	Feature	MC-UG 1D	10.00	10.00	$572.00	$5,720.00
POSTSTACK	Feature	DU 1D	4.00	4.00	$522.00	$2,088.00
PVSEIS3	Feature	DU 1D	4.00	4.00	$533.00	$2,132.00
R03 OpenWorks without Oracle	Feature_Set	ET 1H	93.50	0.00	$298.00	$0.00
R03 PostStack	Feature_Set	MC-UG 1D	4.00	4.00	$750.00	$3,000.00
R03 SeisWorks 3D	Feature_Set	MC-UG 1D	6.00	0.00	$753.00	$0.00
SEIS3D	Feature	MC-UG 1D	6.00	6.00	$692.00	$4,152.00

Figure 7: Invoice for Chargeback

Chargeback is utilizing usage reports to actually create internal chargebacks to departments or groups within the company. There are many different methods to do chargeback. The invoice in Figure 7 represents a method where the department pays a base fee for services and then has a component for overages. Some companies do an actual internal chargeback for everything and make IT a profit center.

Showback is when the same usage reports are utilized to communicate to departments how much their IT usage and requests cost the company without an actual charge to their budget. Utilizing showback reports can really help departments understand IT cost and helps make them more accountable for maintaining good habits.

If storage costs are looked at on a tiered basis for performance, an organization that runs projects should move closed project data to Tier 3 storage instead of Tier 1 storage. It may not even be a question of tiered storage and may just be a good housekeeping process that moves static information off primary SANs or servers.

Allocation Models

Code	Definitions
MC-UG	Maximum concurrent # users/ user group
DU	Distinct # users
ET	Sum of Elapsed Time
Intervals	**Definitions**
1H	Per hour
1D	Per day
1M	Per month
Filter	**Definitions**
5m	Exclude usage shorter than 5 min per instance
15m	Exclude usage shorter than 15 min per instance

Figure 8: Definition of Allocation Models

As mentioned earlier, there are many ways to allocate cost. Figure 8 provides some definitions for application chargeback. Note that robust tools allow for different definitions per application. This is important to map allocation costs to license agreement terms.

The same way different charges for CPU, RAM, or storage may be allocated depending on whether the app is hosted on a local machine vs. cloud, a PPU agreement vs. an enterprise agreement would typically have different allocation models applied.

Cost Center	Application	Type	Elapsed time	UserDays
A4-811278-5A-0-0A-018	goob	Feature	244.21 h	40
A4-811278-5A-0-0A-040	WellPlan	SMS	0.00 h	31
A4-811278-5A-0-0A-045	ArcGIS Desktop	SMS	0.02 h	1
A4-811278-5A-0-0A-062	GeoMatrix	SMS	4.35 h	1
A4-811278-5A-0-0A-111	Google Earth	SMS	3.88 h	8
A4-811278-5A-0-0A-115	GeoMatrix	SMS	1.36 h	6
A4-811278-5A-0-0A-132	OPENcWELLS Datasources	SMS	0.00 h	5
A4-811278-5A-0-0A-140	Pressure Explorer	SMS	3.99 h	2
A4-811278-5A-0-0A-154	GeoMatrix	SMS	0.09 h	2
A4-811278-5A-0-0A-156	OPENcWELLS Datasources	SMS	0.00 h	8
A4-811278-5A-0-0A-159	GeoMatrix	SMS	2.49 h	1
A4-811278-5A-0-0A-234	Pathfinder	SMS	0.01 h	1
A4-811278-5A-0-0A-271	Hysys	SMS	0.22 h	1
A4-811278-5A-0-0A-282	SLM_HYSYS_Process	Feature	0.17 h	1

Figure 9: App Usage by Elapsed Time

Figure 9 is an example of a report that may be used if the license agreement was by elapsed time and the same allocation method is to be used for chargeback.

Application Metering

Feature Name	Max Utilization	Max In Use	Max Available	Elapsed Time
parallel	28.4%	122	453	27,941.1
multiple_realisation	24.8%	51	206	10,402.8
eims_ds_geoframe	1.0%	2	200	120.2
gf_welledit	0.5%	1	204	20.3
dtm_transfer_framework	0.4%	1	260	0.2
gf_geoviz_explore	0.3%	1	304	0.1
gf_lithotoolkit	0.0%	0	304	0.0
eims_ds_openspirit	0.0%	0	200	0.0
gf_ressum	0.0%	0	204	0.0
gf_data_functioning	0.0%	0	204	0.0

Figure 10: Top Ten Max in Use

Figure 10 shows the top 10 max in use by feature including max utilization, max in use, max available and elapsed time. In this example, under the utilization by feature of an app, there are many features that are not being used very much or utilized at all and, therefore, maintenance and support may be cancelled for cost savings.

In addition, the top 2 features never exceed 28.4% utilization or 122 max in use, which could also represent potential savings on maintenance and support. One is licensed for 453 users and the other for 206. These are big deltas that must be analyzed and acted upon.

Maxed Out Licenses

Application Name	Max Utilization	Max In Use	Max Available	Elapsed Time	Denials	Avg Utilization
petrogen3D	100.00%	3	3	349.6		11.15%
struct2D	100.00%	1	1	4.4		1.09%
PSIMENGINE	100.00%	1	1	0.9		0.55%
Editor	100.00%	3	3	191.6		20.96%
OLGAS-3PHASE	100.00%	1	1	23.2		2.26%
GeoStats	100.00%	1	1	0.1	4	0.31%
PIPESIM-NET	100.00%	1	1	238.8		19.30%
gf_flogrid	100.00%	2	2	3.4		0.27%
risk3D	100.00%	1	1	35.6		3.12%
gf_floviz	100.00%	7	7	2,326.5		30.46%
datacheck	100.00%	10	10	18.6		1.13%
office	100.00%	10	10	5,168.5		47.15%
gf_spectral_decomposition	100.00%	1	1	5.6		0.55%
petroflow3D	100.00%	3	3	248.9		7.81%

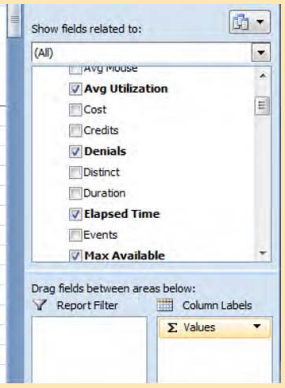

Figure 11: Utilization by Feature

Again, are purchase decisions being made using only one metric? The value of true usage metering is looking at the data from multiple data points to see what the true story is.

In Figure 11, there are feature sets that show 100% peak utilization. Many end users would probably be complaining and request additional licenses. However, look at how often this happens or what the average utilization of those features is. It can be seen that the average in this list is never above 47.15%. Every company has different thresholds, so it becomes a business decision at this point if purchasing additional licenses is warranted or not. However, this report allows decision makers to make an informed decision with all the facts instead of just one data point.

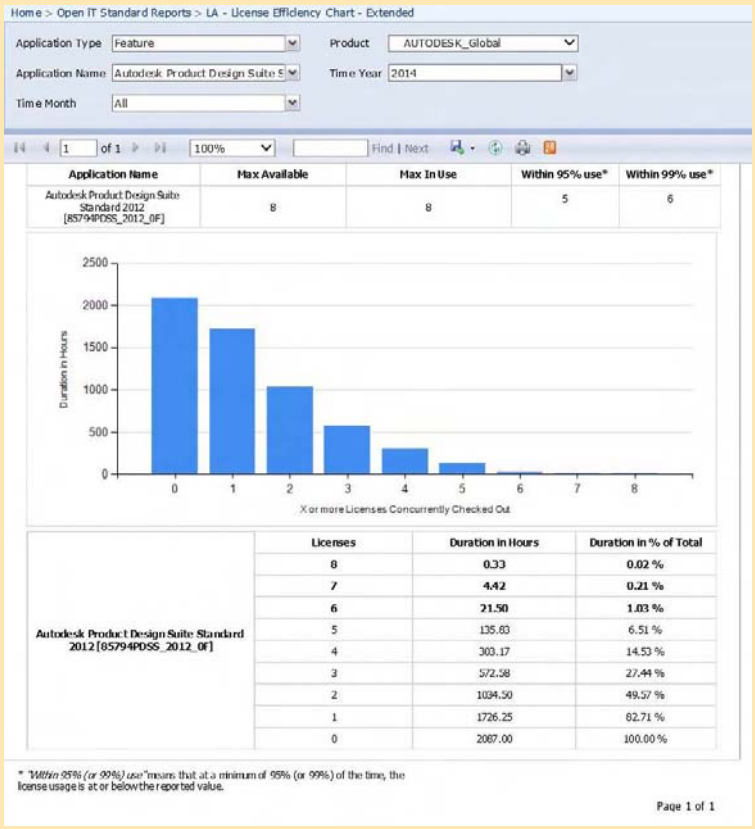

Figure 12: License Efficiency Chart

Figure 12 is another way to look at license efficiency. This chart quickly shows how many licenses are needed to accommodate requests 95% or 99% of the time.

Max Concurrent/Days Since Last Use for: *shou01lmkr21;lgcx, shou01lmkr21;licsrv*
Query Period: *2011-08-01 00:00 to 2011-09-28 00:00*

Application Name	User Id	User Full Name	Max In Use	Elapsed Time	Days Since Last Used
DATALOAD	bap010	James Moore	1	1.17 h	5
	cbw005	Kim Mills	1	56.00 h	15
	cnw009	Russell Billington	1	7.67 h	18
	dam004	Helen Dunham	1	47.50 h	14
	dhs013	Ricardo Corry	1	165.42 h	41
	dkr023	William Bonner	1	0.17 h	39
	gol011	Joan Bryant	1	0.17 h	20
	kbu026	Jeffery Hand	1	0.50 h	18
	ldu015	Adriene Bland	1	2.50 h	41
	nja044	Kimberly Banister	1	22.75 h	1

Figure 13: Named-User Agreements

Named User Agreements

Figure 13 is a report specially designed for administrators to follow up on named user agreements. To optimize license agreements, named users should be reserved for power users, not infrequent users. Named users can be seen by application and what their usage is and also how long it has been since they used the app. This report indicates that these named users are not power users at all and should be converted to a different type of license for better optimization of the license and cost avoidance.

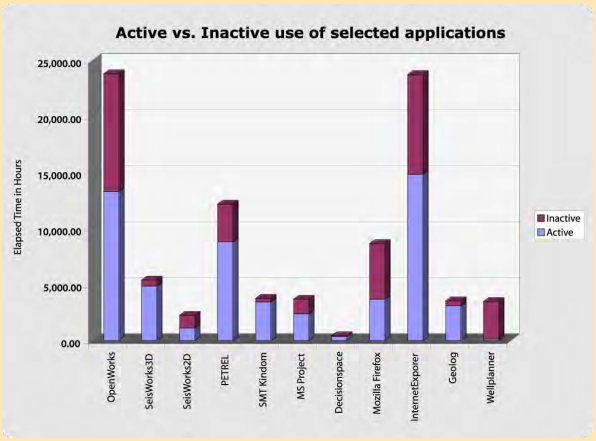

Figure 14: Perceived Use vs Actual Use

The report in Figure 14 can be used to measure not only applications running on a license server, but any and all applications on a workstation or terminal server. This type of data can be used for all types of decisions such as:

- Are people using non-standard apps?
- Are people camping on licenses on license servers?
- Are people using the apps they have requested?

The most common apps that companies want to look at in this category are Microsoft Visio and Project, as well as Adobe.

How many people request a license saying they have to have it to do their job? However, when looking at their usage data, they really only needed it for a short period of time for a specific project. The license is consumed and not being used. It should probably be harvested for use elsewhere.

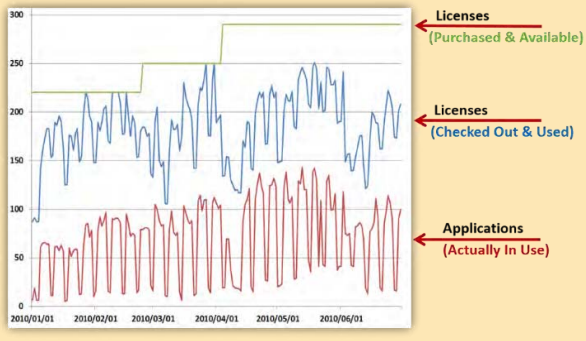

Figure 15: Licenses Checked Out vs. Actually in Use

Figure 15 shows a simple way to view the data report of what has been purchased vs. checked out vs. actually used. This is a good graphic to show management to help them quickly understand the benefits and potential

savings to the company by investing in a metering tool. The green line represents the number of license purchased or available. The blue line represents the license checked in and out from the license server. The red line is the true active usage of the software. That is, how much are the licenses actually being used in the environment. We can see that the customer appears to have purchased more licenses based upon checkouts, but was that really necessary?

What they need to base their decisions on is the red line or the actual usage represented here. If this customer had been looking at these reports, not only could they have saved by cancelling maintenance and support on approximately 100 licenses, but they also would have avoided the cost of purchasing an additional 50 licenses. These are huge savings.

Considering an average industry application costing $30,000, the reduction in maintenance would represent a cost reduction of $600,000, which would more than pay for the tool. However, we can also demonstrate more value. Calculating the cost avoidance of purchasing the additional 60 licenses, that represents a cost avoidance of $1.8M.

Conclusion

Metering, reporting and alerting on cloud usage encourages cost reductions and cost avoidance behaviors. Utilizing showback or chargeback reports keeps departments better informed and responsible for IT resource costs. It changes the conversation between the IT department and their customers, their users and department heads.

Implementing a metering tool is much easier and faster than implementing a complete asset management tool. It also has a very short ROI and, sometimes, can be a good place to start in the overall asset management plan because the cost savings realized by using a metering tool can help pay for additional people to manage IT assets, an IT asset management tool, or other budget items.

A successful cloud management and implementation is planned fully, charged accurately and monitored rigorously. This can be achieved by having a proper usage metering program that delivers the right metrics to help IT and business executives identify and correct under-optimized cloud deployments. Usage metering clearly demonstrates how IT can aid in detecting over-licensing, reduce unnecessary software costs and promote better compliance, best practices and efficient use of IT assets, ensuring a smooth move to the cloud.