



Case Study

The Multi-faceted, Multi-service, Hybrid, Heterogeneous Cloud Environment at Marist College

Executive Summary

When we think of cloud computing architecture, we immediately think: 1) virtualization; and, 2) service delivery models. Cloud architecture is all about maximizing the use of computing resources — and virtualization helps do this by increasing systems/storage/-networking utilization rates. Cloud architecture also makes it possible to deliver computing resources to users in new service delivery forms such as “Platform-as-a-Service” (PaaS), “Infrastructure-as-a-Service” (IaaS), Software-as-a-Service (SaaS). Further, cloud architecture enables the creation of specialized services such as security services, business process management services, business analytics services, complex event processing, and collaborative process services.

What we found when we recently visited Marist College is a highly virtualized, heterogeneous, hybrid (public and private) cloud computing environment that delivers eight distinct services (and dozens of sub-services). Marist currently runs more than a thousand virtual machines within its cloud. And the list of services that it provides to administrators and faculty, students, the local business community, the open source community, and vendors includes:

1. Desktop services for students (e-mail, word processing, spreadsheets and more);
2. A development/test environment for student entrepreneurs and local businesses that need access to PaaS and IaaS to innovate and test new ideas.
3. A dev/test environment for vendors and the open source community (for instance, the development of the System z distribution of Debian Linux is done at Marist);
4. A learning environment that provides thousands of students around the globe with access to their own virtual servers for educational/experimental purposes (Marist runs an open source Sakai environment [a learning management system, research collaboration system and ePortfolio solution] as a cloud service);
5. A “knowledge center” for IBM System z (mainframe) training;
6. A Linux server environment that provides computing resources to the student computing society;
7. A business analytics cloud environment; and,
8. Web hosting services for a variety of associated partners.

In this *Case Study*, *Clabby Analytics* takes a closer look at Marist’s heterogeneous hybrid cloud environment. And we also examine Marist’s relationship with IBM — and its participation in a program known as the “IBM Academic Initiative” (a program focused on helping grow a new crop of mainframe-skilled individuals).

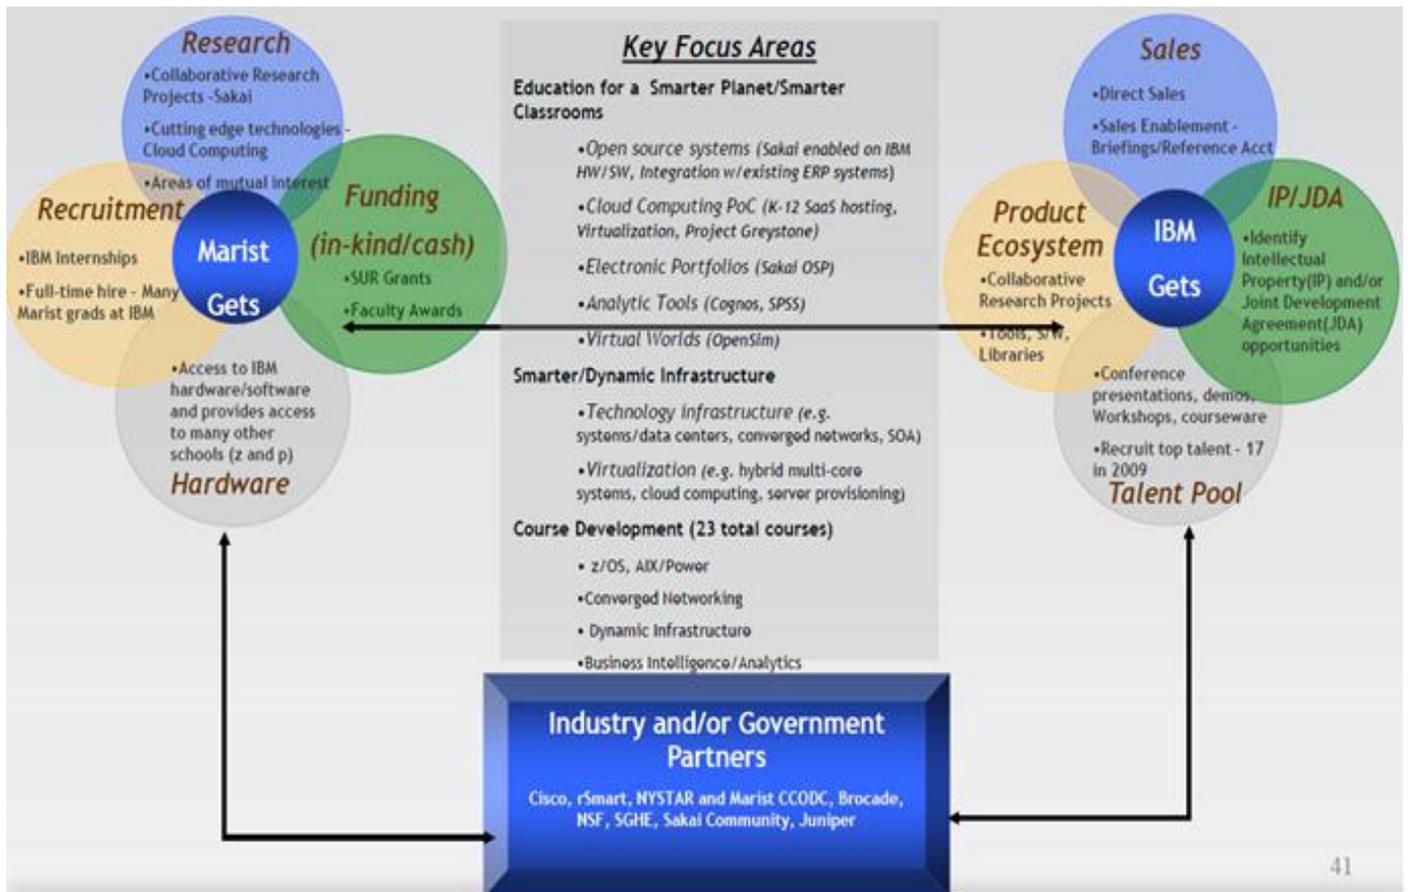
The Multi-faceted, Multi-service, Hybrid, Heterogeneous Cloud Environment at Marist College

Background

Marist College is a private liberal arts college on the east bank of the Hudson River in Poughkeepsie, New York. Its mission is to help students develop the intellect and character required for enlightened, ethical and productive lives in the global community of the 21st century. Marist seeks to distinguish itself by “the manner in which it uses information technology to support teaching, learning and scholarship at both the undergraduate and graduate levels...”

Marist’s location is a huge advantage for the college. Its campus is located within a few miles of IBM's Poughkeepsie research and manufacturing facilities — ideally positioning Marist to serve as an IBM research partner (and also enabling Marist's best-and-brightest students to seek internships and potential employment at IBM after graduation). But the Marist/IBM relationship goes far beyond research and recruitment (see Figure 1). Marist also works hand-in-hand with IBM to provide IT education (particularly in the areas of z/OS, AIX/Power, Linux, converged networking, dynamic infrastructure, and business intelligence/analytics); works with IBM on open source initiatives; and participates in IBM briefings at IBM’s Poughkeepsie visitor center. In return, Marist gets access to IBM hardware, software, and services.

Figure 1 — A Closer Look at the Marist/IBM Relationship



Source: Marist College — June, 2011

The Multi-faceted, Multi-service, Hybrid, Heterogeneous Cloud Environment at Marist College

It is also important to note that Marist is a member of IBM's Academic Initiative — a program designed to promote mainframe education and placement. For the past three years Marist has hosted a conference at its campus where university professors from all over the United States have convened to share mainframe course curriculum ideas and educational approaches with a goal of helping to increase the number of students building mainframe skill sets. (*Clabby Analytics* has attended Marist College's Enterprise Computing Conference [ECC] each of the last three years — and has seen the number of attendees at this conference grow rapidly — and has also seen the number of colleges and universities that support mainframe curriculum grow from about 500 three years ago to more than 900 today). Further, over this timeframe we have seen the initiative move in a new direction: helping mainframe-skilled individuals find jobs. IBM now operates a site called Systemzjobs.com — a website that currently links over 1200 registered, mainframe-skilled individuals with 75 employers.

What Makes the Marist Cloud Different: Heterogeneity and the Number of Cloud Services

Marist College's cloud environment is the third academic cloud that we have examined. Two years ago, *Clabby Analytics* visited the University of Qatar/Texas A&M/Carnegie Mellon cloud site in Doha, Qatar. And, subsequently, we had a discussion with Professor Giuseppe Visaggio about the University of Bari cloud (in Bari, Italy). What we observed as a result of these discussions was that all of these clouds were homogenous implementations — and each offered different services to their constituents. The biggest difference between the Marist cloud and these two cloud environments can be found in the nature and breadth of services offered.

The Qatar Cloud

The Qatar cloud is a homogeneous environment (a large x86 blade environment) designed to teach students basic cloud principles such as virtualization, provisioning, and workload management. This cloud also serves as a research cloud environment where students and professors can access many computing resources in order to run large computing jobs. And, the Qatar cloud offers the local community and the local government access to spare capacity (on a chargeable basis). What we did not see when we visited Doha was a major emphasis on providing a large number of services beyond basic PaaS and IaaS services. This statement is not meant to criticize the Qatar cloud implementation — it merely points out how the three universities use their particular cloud environment.

The University of Bari Cloud

Like the Qatar cloud, the University of Bari focuses on student education — but in this case, this cloud environment is mainframe centric (also making this cloud a homogeneous cloud environment). Students are taught mainframe skills, and are also given access to computing resources to conduct research activities. The Bari cloud differs from the Qatar cloud in that it focuses on providing services to the local business community.

The University of Bari has built three cloud applications to date and expects to build many more in the future.

1. One of the first cloud applications that the university released was designed to make the local fishing industry more efficient. Using this application, fishermen “suppliers” can use a touchscreen device to describe the volume and types of fish

The Multi-faceted, Multi-service, Hybrid, Heterogeneous Cloud Environment at Marist College

being caught. This information is then placed in a database that is made available to buyers who can bid on the catch. Upon completion of a successful bid, fishermen can start to box their catch, enabling their fish to be shipped more rapidly to buyers. This application streamlines the selling process as well as the delivery process.

2. The University of Bari's wine market application behaves in very much the same way as the fishing application — matching wine producers to wine resellers. And again, the benefits are similar in that the selling process is streamlined, and stock is delivered to where it is needed more efficiently.
3. The University has also developed an application that it calls “MoniCA (Monitoraggio e Controllo Adattivo)” – a logistics application that monitors truck activity in real time. This application collects data related to the temperature, humidity, and shocks delivered to the cargo within the truck — helping those who use this application to avoid damaging the goods being delivered, as well as to avoid rough roads on delivery routes. This type of monitoring/data collection and database analysis activity is what IBM frequently refers to as a “smarter planet” application.

Notice that the university's applications are delivered as services to the consuming public. The backend architecture is completely transparent to the user — all that the user sees is an interface to the backend. This is one reason why cloud computing is becoming very popular with end-users.

The Marist Cloud

What makes the Marist cloud different when compared to these other two cloud environments is that it is a heterogeneous cloud environment — and that it offers many services to a much broader audience.

The Marist cloud systems environment consists of:

- 2 IBM System z mainframes (running z/OS, z/VM, and Linux operating environments);
- 11 IBM Power Systems servers (running AIX [Unix] and Linux);
- 75 x86-based servers (running Windows and Linux); and,
- 70 Terabytes of state-of-the-art IBM storage.

This environment is highly virtualized (the college runs 955 virtual servers on its System z mainframes; 40 on its Power Systems servers; and more than 50 virtual machines on its System x [x86] blades).

As stated in the Executive Summary, the Marist cloud offers eight services — some of which are similar to those offered by the Qatar and Bari clouds (such as student services and access to virtualized research servers) — but many that are distinctly different from the Qatar and Bari clouds (including an open source collaboration environment; major research projects with IBM; a System z knowledge/training service; and business analytics research and services). The next section explores these Marist cloud services offerings in greater detail.

The Multi-faceted, Multi-service, Hybrid, Heterogeneous Cloud Environment at Marist College

Marist Cloud Services

Marist offers eight distinct services that to its constituents through its cloud architecture. In this section we examine these services more closely.

Desktop Services for Students

Many colleges and universities offer their students access to word processing, presentation, spreadsheets and other “office” applications — and Marist is no exception. Marist provides its students with the software image that contains the office software, an e-mail environment (Foxmail), and large file drop boxes that can be used for file sharing. The Marist cloud overlays security functions across these environments.

Dev/Test Environment for Students

All three cloud environments described in this work provide students with access to virtual servers for development and testing purposes. What made Marist stand out was the amount of cloud processing power that it offers its students. Currently, Marist students have access to more than 600 virtual servers for dev/test and learning activities.

What was also particularly noteworthy about the student dev/test environment is that Marist makes computing power available to student entrepreneurs (and local businesses with Marist ties — including some alumni) in order to experiment with business models and ventures.

Dev/Test Environments for Vendors and the Open Source Community

We separate dev/test for students from dev/test for vendors/open source because the nature of the work being conducted in a vendors/open source dev test environment is distinctly different from the nature of the work conducted by students.

When working with vendors, colleges often work under a grant arrangement where funds, or goods, or services are exchanged for research conducted by the school — and, accordingly, a lot more rigor must be applied to the cloud services being delivered.. In the case of student dev/test, students can build systems images, experiment with those images — and should a crash occur, students can simply build a new server and start over. When working with vendors, crashing systems to learn by trial-and-error may not be a desirable dev/test approach.

To us, we were surprised to find how much vendor dev/test takes place at Marist — and we were pleased to see the breadth of products that Marist students were working on. Some of the vendor dev/test environments include:

1. The integration of IBM’s WebSphere Application Server and DB/2 database with open source Sakai collaboration software;
2. Rational software (an IBM application development environment) dev/test;
3. SPSS and Cognos (IBM business analytics software) dev/test; and,
4. Linux on z (testing and proof-of-concept activities).

The Multi-faceted, Multi-service, Hybrid, Heterogeneous Cloud Environment at Marist College

Open source dev/test efforts include open source Sakai work; work on System z for the Linux Foundation; work on Slackware — and work on Debian Linux. We see the Sakai work as a distinct cloud service offering.

The Sakai Environment

Marist efforts in the Sakai space go far beyond dev/test. Marist is very actively involved in promoting Sakai as a learning management system, research collaboration system and ePortfolio solution — and makes Sakai services available on its cloud and encourages the Sakai community to make use of those services. (As an example, Marist hosts quality assurance servers for the Sakai community).

It should also be noticed that Marist work in Sakai — and its recommendations — led to the adoption of Sakai by the State Bank of India (approximately 300,000 users!).

The Knowledge Center

Marist specializes in mainframe education — and has built a mainframe educational cloud service that calls its “knowledge center”. This public cloud offering provides access for thousands of students (including Marist students as well as domestic and international students) to mainframe course curricula.

Using the knowledge center, students and professors can access z/OS images from around the globe for learning activities. From this z/OS starting point, students can launch a DB/2 database; learn the COBOL and PL/1 programming languages; learn about IBM's CICS and Web sphere environments; launch a variety of third-party applications including MVS Quickref from ChicagoSoft, Easytrieve from CA, z/XDC Debugger from Cole Software, XMITIP from LBDS Software, SELCOPY Productivity Tool from CBL, and PathPoint Development Tools from PathPoint Software.

This cloud service is heavily used. Since October, 2010, over 214 courses have been taught using knowledge center cloud services. Approximately 500 professors have made use of the services. Over 200 schools worldwide have participated in knowledge center activities — in schools that include the United States, Lithuania, and all the way to Queensland, Australia. Incidentally, ratings have been high and feedback positive from both professors and students for the quality of service offered by this Marist cloud service.

Linux for the Student Computing Society

Marist has created a specific cloud computing environment for members of the student computing society. Members of this society have access to Linux on System z, and can thus develop Linux on z skill sets (which are hotly desirable in today's IT marketplace).

To help develop these skills, Marist provides the students with access to Linux services on its primary mainframe (a z10 Business Class server) that runs two general processors and offers eight Integrated Facility for Linux (IFLs — Linux specialty processors) and up to 80G memory. These students are also given access to two LPARS (logical partitions) that run IBM's z/VM operating environment (so they can develop z/VM skills as well).

The Multi-faceted, Multi-service, Hybrid, Heterogeneous Cloud Environment at Marist College

Two Business Analytics Cloud Environments

Business analytics is a hot market segment in the IT field — and Marist is addressing this segment by offering two business analytics/business intelligence cloud services: 1) a cloud service for administrators; and, 2) a cloud service for academic use. These services make use of IBM SPSS and Cognos software — and can be used in a virtual computing lab setting to teach business intelligence course on-line.

Web Hosting for a Number of Student Organizations and College Partners

In addition to hosting dev/test environments and operating knowledge centers, Marist also provides web hosting cloud services for various student organizations as well as college partners (including the FDR Presidential Library, Hudson River Valley Institute (HRVI), Poughkeepsie Town Police, Hudson River Environmental Society, Institute for Data Center Professionals (IDCP), Center for Collaborative and On Demand Computing, Center for Teaching Excellence, the New York Media Experience Program, and the Enterprise Computing Community (ECC). And the sheer volume of this activity is quite impressive (as shown in Figure 2).

Figure 2 — Marist Web Hosting Activity

Metrics	FY 2009	FY 2010	FY 2011
Pages Viewed	15,830,000	15,952,000	16,176,000 (10M External)
Per Month Average	1,300,000	1,330,000	1,348,000
Unique Visitors	1,386,000	1,145,000	1,182,000 (1M External)
Per Month Average	115,500	95,500	98,500
New Visitors	21%	21%	21.4%
Functional Areas Published On the Site	230	280	290
Average Monthly Updates Made to the Site	300	600	710
Staff Members Who Actively Post	50	68	73

Source: Marist College — June, 2011

Summary Observations

This case study focused primarily on the multi-faceted/multi-service cloud environment at Marist. But it should not be overlooked that the Marist IT department also drives a \$130 million business (the college itself). It provides access by the staff to advancement, executive self-service, finance, financial aid, general, human resources, student, and enrollment management systems. And it provides access by students to college services, desktop services (including e-mail, document management, and spreadsheets), as well as access to large collaborative environments. We also must note Marist runs SunGard's higher education Banner software portfolio on its mainframe — Marist worked with SunGard to help port this software and told us that the SunGard developers were surprised at how straightforward the effort was to port to Linux on a mainframe.

The Multi-faceted, Multi-service, Hybrid, Heterogeneous Cloud Environment at Marist College

This case study merely touched the surface of cloud activities taking place at Marist College. For instance, we did not go into depth on the Rational application development toolset and how it is being used in a cloud environment. Nor did we talk at length about the on-campus gigabit backbone, nor the security environment and firewalls that protect Marist clouds. And we did not have room in this report to expand on how Marist's Cognos cloud offering provides educational services for Universidad de Buenos Aires and Universidad de Belgrano, Argentina. There is so much cloud activity going on at Marist that it is simply mind-boggling.

What did describe is a large, heterogeneous, hybrid computing environment that makes heavy use of virtualization technology to get maximum computing capacity usage out of the servers that Marist uses. And we described a wealth of services provided by Marist to its students, faculty, and administrators — as well as to enterprises, vendors, and the open source community.

It is this wealth of services that really caught our eye. We've seen other academic cloud environments that are primarily focused on using clouds for learning activities, for dev/test, and to serve the local community — but we've never seen cloud environment with so many services being delivered to constituents. And the beauty of this Marist cloud is that Marist will continue to add new and innovative services as students, teachers, and vendors ask for more and more access to computing resources.

Given that *Clabby Analytics* is a huge believer in mainframe architecture, we must also note that Marist's clouds make heavy use of IBM's System z to provide much of the horsepower that drives its cloud services. We must also note that the mainframe curriculum that Marist offers — and the ready access to mainframe tools and utilities — is going a long way toward building the next generation of mainframe managers and administrators.

And finally, we'll end on a philosophical note. When we asked if building the Marist cloud was a daunting job, one of our hosts responded "hey, it mostly runs on the mainframes. And mainframes have offered all of the tools and utilities — and the advanced virtualization facilities — needed to build clouds for decades. Building a mainframe cloud is no big deal"...

Clabby Analytics
<http://www.clabbyanalytics.com>
Telephone: 001 (207) 846-6662

© 2011 Clabby Analytics
All rights reserved
July, 2011

Clabby Analytics is an independent technology research and analysis organization. Unlike many other research firms, we advocate certain positions — and encourage our readers to find counter opinions — then balance both points-of-view in order to decide on a course of action. Other research and analysis conducted by Clabby Analytics can be found at: www.ClabbyAnalytics.com.