


Reshaping Applications and Business Intelligence at UTS

David O'Connor, Susan Gibson





- Complex architecture - big mix of legacy, modern, bespoke etc
- Pace of technology change
- Facing demands to be more agile
- Cost – do more with less/same
- Governance, security, privacy
- Many vendors to manage
- Explosion of data
- Managing identity across entire education lifecycle

To support the UTS vision of being a world leading university of technology – IT need to adapt.

For context - 80 staff responsible for everything from BI and DWH, ERP, learning systems, web, intranet, mobile, dev and integration etc.

- 2 year strategy to change the application landscape
- **Our mission**
 - Create a challenging, creative, professional, and fun team environment...
 - ...that enables us to consistently deliver innovative, user-centric, and appropriately scoped technology solutions...
 - ...that help drive UTS towards its goal of being a world leading university of technology.

1. Strategic relationship management
 2. Application technology
 3. Project management
 4. Operational strategy
 5. Staff learning, development & workplans
 6. Modernisation & continuous improvement
 7. Contractor & strategic
- 11. Be a productive part of the UTS community**
- Everyone contributes to the UTS community in some way
 - We get out and talk about what we do
 - We consider a 'UTS-first' policy for expertise
 - We employ UTS students in our teams wherever possible
11. Be a productive part of the UTS community
 12. Supporting an innovative culture

Supporting an Innovative Culture – Specifically in Data Warehousing and Business Intelligence

- Students
- Courses
- Enrolments
- Research Income
- Research Publications
- Learning Management
- HR
- Financials
- Light levels
- Room Temperature
- Humidity
- Air quality
- Energy
- Water Use
- People Counters
- Data collected in real time

We need a data warehouse ecosystem that can cope with:

- Hundreds of source systems
- Large scale data
- Data in all flavours
 - Structured
 - Unstructured
 - Real-time

Needs to be highly scalable

- Many different options and many price points
- Previously database servers were not used for heavy analytics but data would be transferred to a middle tier (e.g. cube)
- Led to compromise – could not drill through edge of cube- create new cubes – cube explosion – reconciliation and administrative nightmare
- Databases have developed technologies specific to data warehousing
 - high performance
 - large scale data sets

Multiple on premise options – require high capex

Recently Amazon Web Services released a new service globally – made available in Australia March 2014

- Columnar data store, data compression, parallel queries across multiple nodes - very high performance
- Elastic cloud architecture – can scale as required
- Cost are actually very low – no capex investment required.
- SAAS
- Increased development agility
- Petabyte scale data warehouse

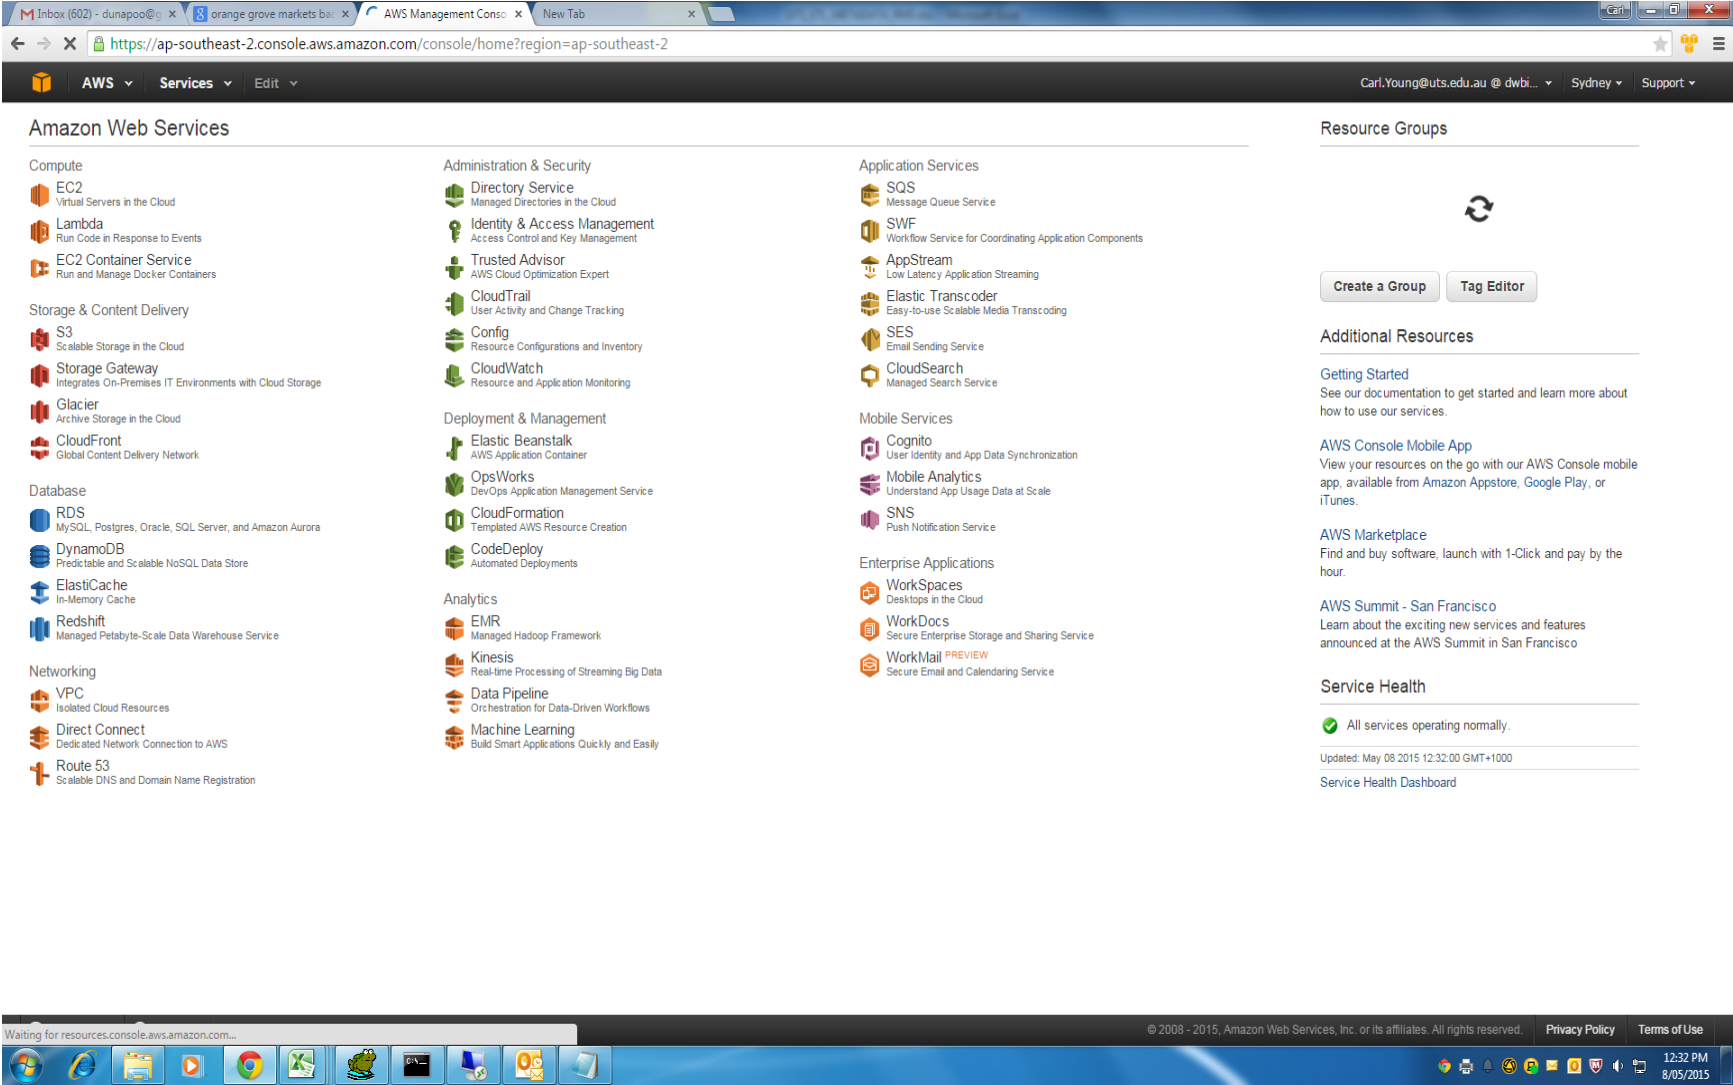
Data in Data Warehouse

- Course and Subject Details
- Student Lifecycle data including enrolments, attrition, completion
- Fee Income
- Results and Grades
- UAC
- Research Data

- Biggest Table – 2.2 Billion rows (daily snapshot of student load for 18 months)
- Sub 10 second queries to summarise load by day
- 813 entities

Cost

- 1 TB Development and Production Environment
- Combined \$5,000 per month



The screenshot displays the AWS Management Console interface. At the top, the browser address bar shows the URL <https://ap-southeast-2.console.aws.amazon.com/console/home?region=ap-southeast-2>. The console header includes the AWS logo, navigation menus for Services and Edit, and user information for Carl.Young@uts.edu.au.

The main content area is titled "Amazon Web Services" and is organized into several columns of service categories:

- Compute:** EC2 (Virtual Servers in the Cloud), Lambda (Run Code in Response to Events), EC2 Container Service (Run and Manage Docker Containers).
- Storage & Content Delivery:** S3 (Scalable Storage in the Cloud), Storage Gateway (Integrates On-Premises IT Environments with Cloud Storage), Glacier (Archive Storage in the Cloud), CloudFront (Global Content Delivery Network).
- Database:** RDS (MySQL, Postgres, Oracle, SQL Server, and Amazon Aurora), DynamoDB (Predictable and Scalable NoSQL Data Store), ElastiCache (In-Memory Cache), Redshift (Managed Petabyte-Scale Data Warehouse Service).
- Networking:** VPC (Isolated Cloud Resources), Direct Connect (Dedicated Network Connection to AWS), Route 53 (Scalable DNS and Domain Name Registration).
- Administration & Security:** Directory Service (Managed Directories in the Cloud), Identity & Access Management (Access Control and Key Management), Trusted Advisor (AWS Cloud Optimization Expert), CloudTrail (User Activity and Change Tracking), Config (Resource Configurations and Inventory), CloudWatch (Resource and Application Monitoring).
- Deployment & Management:** Elastic Beanstalk (AWS Application Container), OpsWorks (DevOps Application Management Service), CloudFormation (Templated AWS Resource Creation), CodeDeploy (Automated Deployments).
- Analytics:** EMR (Managed Hadoop Framework), Kinesis (Real-time Processing of Streaming Big Data), Data Pipeline (Orchestration for Data-Driven Workflows), Machine Learning (Build Smart Applications Quickly and Easily).
- Application Services:** SQS (Message Queue Service), SWF (Workflow Service for Coordinating Application Components), AppStream (Low Latency Application Streaming), Elastic Transcoder (Easy-to-use Scalable Media Transcoding), SES (Email Sending Service), CloudSearch (Managed Search Service).
- Mobile Services:** Cognito (User Identity and App Data Synchronization), Mobile Analytics (Understand App Usage Data at Scale), SNS (Push Notification Service).
- Enterprise Applications:** WorkSpaces (Desktops in the Cloud), WorkDocs (Secure Enterprise Storage and Sharing Service), WorkMail (Secure Email and Calendaring Service).

On the right side of the console, there are sections for "Resource Groups" (with "Create a Group" and "Tag Editor" buttons), "Additional Resources" (including "Getting Started" documentation, "AWS Console Mobile App", "AWS Marketplace", "AWS Summit - San Francisco", and "WorkDocs"), and "Service Health" (indicating "All services operating normally" as of May 08 2015 12:32:00 GMT+1000).

The bottom of the screenshot shows the Windows taskbar with various application icons and the system tray displaying the time as 12:32 PM on 8/05/2015.

- Finding the right people and keeping them
- Legal, Governance, Security, Architecture
- Simple technical issues such as ODBC, JDBC issues, getting out of the UTS network, ensuring VPC is secure, forcing queries back to database
- Consider how long you can stay in a situation where you are making compromises
- None of the challenges were insurmountable but you need to be persistent