

THE REAL GREAT-GRANDFATHER ALBERT

A new wave of enthusiasm has recently swept across our household - we've become intensely interested in genealogy! I must confess that I had never really felt motivated to trace my ancestors, but once you start it becomes curiously addictive. It all began when my daughter Becky found her way onto the Genesreunited.com website, an invaluable resource for anyone with family links in the UK. The site has tens of thousands of members who build their family trees online and share the results of their research. The great thing is that each member can ask to be alerted automatically if specific names or dates are added by other participants, and they can then get in touch and compare notes on their respective forebears.

Now, I admit that the whole process is a heck of a lot easier if your surname is Lillycrop. There are only 165 instances of that name on the Genesreunited database, and most are probably related in one way or another. If our name were Smith or Brown we might have found those ancestors a lot harder to trace.

We have made some interesting contacts through that website. Becky was thrilled to receive a message from a woman who turned out to be a third cousin, and who had amassed a treasure-trove of information about her (and my) great-grandfather. She initially mailed us because she was skeptical about some of the names and dates on our family tree, many of which were based on my father's earliest recollections and were little more than informed guesswork. The exchange went something like this:

"Are you sure your Albert was born in Devon in 1870? I have London as his place of birth, and in 1874. I'm sure his father moved from Devon to London in 1852."

"Oh, that's strange. Are we talking about the same Albert?"

"Yes, I've checked the 1901 census, and there was only one Albert Lillycrop who remotely fitted the description."

"Thanks - your information sounds much more reliable than ours. We will update our family tree."

It was indeed the same Albert, but by describing him in different ways we had erroneously created two different people. With the benefit of shared information, the two Alberts became one.

It's just the same in the average data center. Part of the reason for the complexity of system management is the sheer diversity of hardware and software assets and the way that they are identified. Monitoring tools can only work effectively if they have the necessary information to present a clear, complete picture of all the hardware, software and network assets within the enterprise. Similarly, for service management to work smoothly, we need an end-to-end change management process that identifies and records all changes to the relationships between applications, data items, and physical devices. Often we have

conflicting, duplicated, incomplete or overlapping views of IT assets which - like our family trees - give us multiple versions of the truth.

In today's SOA-dominated world, where services are discovered and integrated on the fly, the need for a single central source of reference is more important than ever. That's where the Configuration Management Database (CMDB) comes in, an essential focus for all change and asset management activity. Companies that are building their service management strategy around ITIL or similar 'best practice' methodologies are increasingly seeking ways to create a reliable CMDB at the heart of their IT environment.

Of course, in a heterogeneous world and with multi-vendor tooling, there might be several (partly incompatible) CMDB implementations involved, particularly if application traffic and data transfer cross company boundaries. So it's encouraging to see some of the heavyweights of the systems management world collaborating on a new standard to overcome this problem. BMC, IBM, Fujitsu, HP and CA are working on what's described as "an open, industry-wide specification for sharing information between Configuration Management Databases and other data repositories." The standard will allow details of servers, storage devices, networks, middleware, applications and data to be federated and shared across heterogeneous environments, promising to help users to gain a far clearer view of changes within their IT environment.

We see many standards come and go in this business, but this working group is one to be encouraged. Unless users have a consistent view of their IT assets and the relationships between them, they cannot be sure that they have a single version of the truth - and that requirement is fundamental to improving service management.

Maybe it doesn't really matter whether Albert was born in London or in Devon (fascinating though it is to find out!), but some truths can have far deeper repercussions.