2010-11 Unix Vendor Preference Survey

GCG's Server Vendor Preference Surveys are conducted with one goal: to cut through the marketing spin, conflicting statistics, and half-baked industry punditry and find out what *real* enterprise customers *really* think about the challenges they face, the equipment they use, and the major vendors that manufacture and support it.

Each GCG survey respondent actually works in a data center and, as you'll see in the demographics that follow, has the depth and breadth of experience necessary to provide keen insight to what's happening on the raised floor.

This is the fifth edition of our Unix survey, and it's quite comprehensive. This survey was in the field from 4Q 2010 through 1Q 2011. In addition to touching on the topics we're covering (below), this document also contains the demographics of this survey, an explanation of scoring and survey methodology, and ways to get more information about GCG's research offerings.

Partial List of Survey Areas

Watch for our upcoming releases of data on a wealth of topics, including:

- **Unix Usage:** Are Unix-based platforms still relevant? Are they strategic? Do customers expect to use more Unix over time?
- **Virtualization/Clouds:** To what extent have customers virtualized their Unix systems? What benefits are they seeing? Are they running to the cloud? We see some troubling trends beginning to emerge.
- **Facilities:** How big a problem is power/cooling/floor space? Is it getting better or worse?
- **Commercial Unix & Linux:** Do customers see Linux as a substitute for their current commercial Unix workloads?
- **Buying Criteria/Brand Differentiation:** What are the most important factors customers weigh when deciding between different Unix brands? Do they think the brands are significantly different, or pretty much the same?
- **Vendor Face-Off:** Customers directly compare the major Unix server vendors (Hewlett-Packard, IBM and Oracle) on a wide range of technical and support criteria, including:
  - System performance, availability, RAS, technical quality, fit & finish, sales & delivery
  - System manageability and management suites
  - Vendor services & support, R&D, system design, roadmaps, etc.
  - Innovation & data center value-add
- **Installed Base Deep Dive:** Each customer identifies which major Unix vendor (HP, IBM or Oracle) is most prevalent in his/her data center. Then we ask: Why did your organization choose that vendor? Are you sticking with that choice? What are your future purchasing plans?

It's a lot of stuff... and clients receive much more, including in-depth analysis, historical data, and custom data cuts, i.e. SMB vs. large enterprise.

2010-11 Unix Server Vendor Preference Survey Demographics:

Organization Size - Total Employees

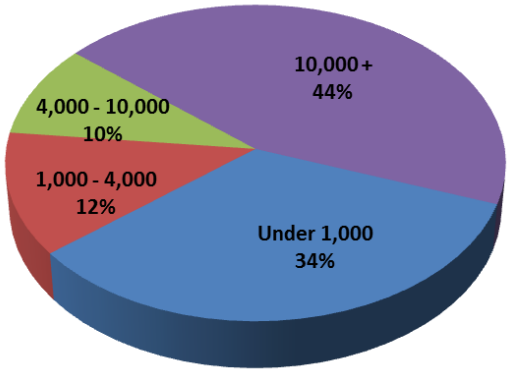

Our surveys are aimed at actual data center workers. We believe that they have the best handle on how their organizations use technology in general as well as specific hardware, software, and services. They have multi-product and multi-vendor experience, and they understand what their businesses and IT operations are planning for the future.

The total number of respondents to this survey was 306, with 54% of respondents in mid-size and large organizations of 4,000 employees and above. The rest are employed by organizations with fewer than 4,000 total employees. Every organization uses at least some commercial Unix in their IT infrastructure.

Physical Unix Servers

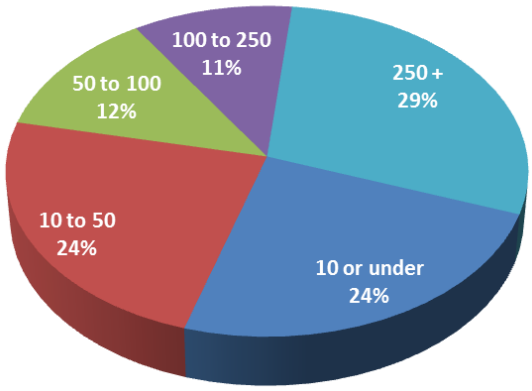

Each respondent was asked how many physical Unix systems they either directly controlled or had intimate knowledge about. As can be seen from the chart, it's quite a range; 30% of the respondent base manages 250 or more physical systems, and 52% manage 50 or more Unix servers. It's important to note that these are physical systems – with virtualization, these platforms are running many more workloads than the system count would imply.

Current Unix Brands in Data Center

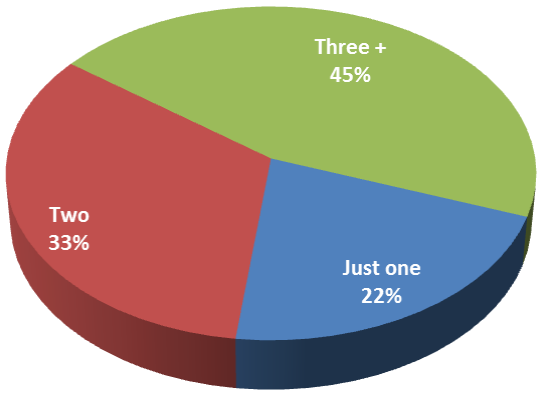

The vast majority of our respondents (78%) are currently using two or more Unix brands, with almost half running all three major Unix operating environments. There's been little change in this pattern over the five years we've been doing this survey.

This is despite the best efforts of vendors to encourage customers to standardize on one commercial Unix to the exclusion of all others. They've all launched various 'take out' campaigns and trade-in programs with the goal of eliminating competitive platforms. While there may be some localized success with these

efforts, we don't see much evidence to support the idea that customers are moving toward commercial Unix standardization.

In addition to currently using more than one brand of Unix, our respondents have had significant experience with every major brand. This range of experience with various brands, along with experience in differently-sized organizations, is what makes our survey respondents such a knowledgeable crowd and gives their opinions such validity. They're smart guys and gals; they know what's happening on the data center floor, and they know where the business is going. That's why we value their opinions so highly.

Survey Scoring & Methodology

In the "Vendor Face-Off" results you'll see something called a VPI Score used as our scoring metric. The easiest way to explain our VPI scoring metric is to lay out a simple example. Let's say this is our annual Unix Vendor Preference survey, and we had 1000 respondents: five hundred of them have standardized on Oracle, two hundred are strong HP customers, and three hundred chose IBM as their dominant Unix vendor. When asked which Unix vendor has the best-dressed salespeople, four hundred participants responded Oracle, three hundred picked IBM, and three hundred said that HP salespeople were particularly natty dressers.

Which Unix Vendor Has the Best-Dressed Salespeople?

Company	Number of Respondents	Number of Votes Received	VPI Formula	Normalized Score
Oracle	500	400	400/500 x 100=	**80**
IBM	300	300	300/300 x 100=	**100**
HP	200	300	300/200 x 100=	**150**

While the raw scores favor Oracle, the normalized score (which is simply the number of "votes" divided by the number of respondents x 100) shows that HP is the winner of this beauty contest. Why? HP captured more first-place votes than the number of respondents in its self-selected base. In this simple example, 100 Oracle customers defected and voted for HP. These are the kinds of situations we want to uncover with our surveys.

We call this normalized number the **Vendor Preference Index (VPI)**. The VPI computation yields an easy-to-understand score for each vendor and a gauge of installed base loyalty. A VPI score greater than 100 means that the vendor in question was selected by a number of respondents greater than the number who have standardized on that particular brand of server. VPI scores greater than 100 are very good.

A VPI score of exactly 100 means that the vendor was chosen as a leader by exactly the same number of respondents as those who have standardized on that vendor. VPIs of less than 100 are, of

course, not so good, and mean that the vendor in question has suffered defections (at least in terms of survey voting) from their own installed base.

While there are certainly more complicated ways to compute the results of a survey such as this, we believe that this method captures the data we are looking for: a fair comparison of how customers perceive vendors. We're trying to understand why customers buy what they buy, and what might prompt them to change their minds about systems and vendors.

The survey itself uses a web-based survey engine hosted on our corporate website. Our survey reaches two main groups of participants. One group is our "qualified" respondents who have taken previous surveys – we've already confirmed that they're data center managers, architects, technicians, etc. The other group is comprised of readers of *The Register* and other technical publications. We've found that this mix gives us a very solid base of knowledgeable respondents.

The survey has several mechanisms in place to prevent respondents from 'ballot box stuffing' by taking the survey more than once. We also have a token-based access mechanism that ensures that only those respondents we invite can take a particular survey.

Much of our effort is focused on data cleaning. This includes testing responses to make sure they are valid, following up with respondents to ask for additional clarification, and making sure our final data set is comprised of knowledgeable IT professionals who are giving us their honest and thoughtful opinions.

GCG Surveys & Offerings

Our publicly-released results are just a glimpse into the massive data set generated by our surveys. GCG uses these findings (along with data from other sources, client feedback, and other activities) as the foundation for our many research and consulting offerings.

We serve our clients in a number of ways, including strategy and product consulting, custom research, and brand evaluations. Clients can obtain access to historical data and customized survey results (different cuts of the data to highlight, for example, SMB respondents or those who are heavily virtualized). We also provide clients with in-depth GCG analysis delivered via research reports that highlight survey data in the context of customer needs or vendor offerings. For further details or to find out how to get access to these additional offerings, please contact us via email or phone.

Gabriel
CONSULTING GROUP
phone / 503.372.9389
gcginfo@gabrielconsultinggroup.com
www.gabrielconsultinggroup.com