DATA CENTER

Data center terminology that will get you hired

Meredith Courtemanche

A photostory

Define the perfect data center job interview

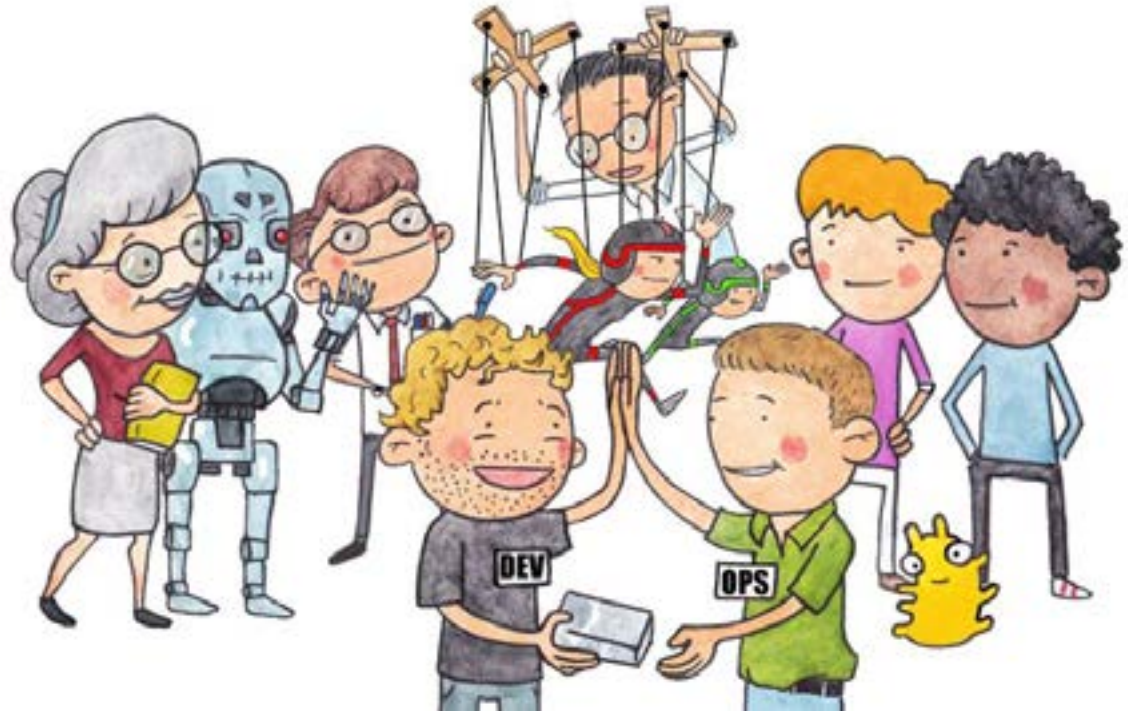

BUSINESSES NEED TO know their back-end IT infrastructure is ready to move the company forward.

Modern IT organizations work closely with an enterprise's business leaders to create, scale up and tear down data center assets at a rapid pace. Flexible and agile IT isn't just for the Web-scale data center operators like Facebook and Google. Enterprises -- from steel companies to coffee purveyors to insurance providers -- are improving IT operations and realigning data centers for an era when digital enables business.

To land a data center job in this evolving space, work relevant, change-focused terminology into your interviews for engineer, operations or manager positions. Simply rattling off DevOps and automation terms won't impress an interviewer, nor will a prevaricating job candidate who lacks the fundamentals to grow in a data center position.

We've created fun visual cues in this slideshow of hot data center terminology to help you understand and describe modern data center operations. Establish your knowledge of IT infrastructure scalability and resiliency, culture and business trends as well as other defining developments while leaving a strong impression on your future employer.

For more preparation before the big interview, check out resources at the end of the slideshow on sample DevOps, sys admin and other data center job interview questions.

Michael Kloran

Define change management as a way to control data centers from above

CHANGE MANAGEMENT GIVES a single administrator control over the actions and performance of an entire fleet of servers.

We define change management as a systematic approach to support the sometimes unpredictable evolution of IT workloads and technologies in the data center. It is about automating workflows, configuration changes, patches and updates.

The definition of change management is increasingly merging with terms for data center automation and orchestration. Configuration and change management tools from vendors as diverse as SolarWinds and IBM give fewer admins control over a greater number of servers, network nodes and storage devices.

When the interviewer asks how you manage your time, or how you would improve time management in the data center, bring up change management. It's also an integral part of successful business project launches and expansions. If your data center can't adapt, your business only has two choices: abandon its growth projects or seek the back-end infrastructure support elsewhere, as in a public cloud and shadow IT.

Michael Kloran

Define continuous integration and delivery as Dev meets Ops

EVERYONE HAS THE same goal in a good IT organization: Put out the best applications and services on stable, sustainable infrastructure. DevOps helps keep that goal on track with continuous integration and delivery.

While developers focus on what's new and next, operations prefers stability and consistency.DevOps is a methodology or culture of IT wherein developers and operations teams frequently communicate about and collaborate on application and service releases. Many modern IT shops expect systems administrators, developers and infrastructure managers to implement continuous integration and continuous delivery.

Continuous integration means that iterative software changes are immediately tested and added to the larger code base. If a change creates issues, such as runaway storage use, that problem can be isolated and rolled back quickly.

Continuous delivery takes committed changes to the code and shepherds them through peer review and user acceptance testing. It is an extension of continuous integration. While many DevOps tools are designed to bring code through both CI and CD, in practice developers and operations teams use different tools in different ways.

Another CD term, continuous deployment, is a goal for DevOps organizations, though often fraught with real-world complications. Continuous delivery means that changes to the application or service are ready to be deployed because they are tested in a production-like environment. Meanwhile, continuous deployment means these changes go straight to production without awaiting human intervention.

Michael Kloran

Define data center resiliency as withstanding destruction

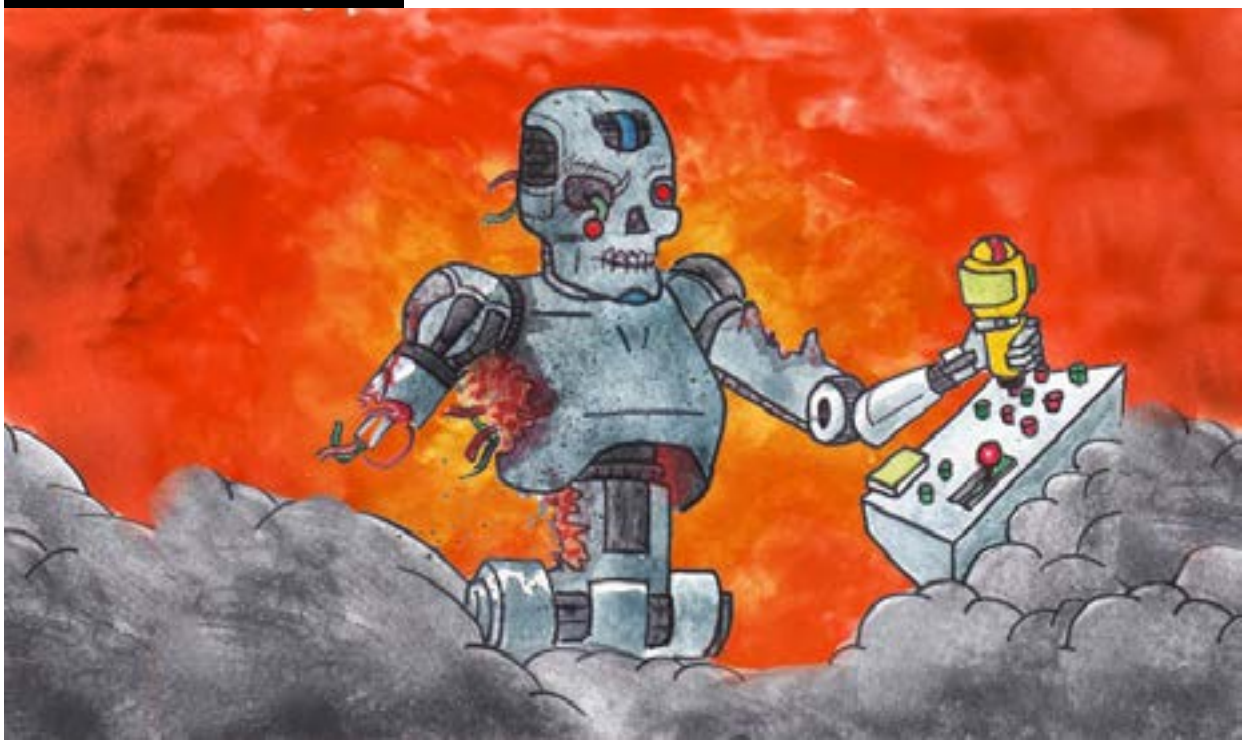

RESILIENT DATA CENTERS withstand damage to the network cabling, power outages, unexpected server downtime, attacks or surges in user demand and other events by avoiding single points of failure and using adaptable IT systems.

In terms of physical infrastructure, data center resiliency increases by installing multiples. If one PDU fails, another is at the ready to take over its workload. If one carrier goes down, the business can run all its workloads over a second carrier's connections. N+1 redundancy means that there's one extra piece of equipment for however many pieces you have operating. N+M could mean two, three or any number of multiples for that equipment. 2N designated an infrastructure that's completed duplicated for protection.

The definition of data center resiliency has evolved as IT technologies such as virtualization and micro services joined traditional uptime boosters like N+1 redundant power supplies and equipment. Be prepared to discuss distributed applications and server clusters in a job interview.

Distributed applications run on multiple servers simultaneously, and may be architected as independently deployed micro services. Server clusters prevent outages with high availability or fault tolerant virtualization, allowing VMs to jump from one server to another in the event of a problem.

Some data centers today are designed to fail, meaning that workloads easily jump off of problem servers and onto new ones, across a server cluster and even across multiple data centers. This architecture is moving from the domain of Web-scale content providers into enterprise IT.

When you talk about data center resiliency to prospective employers, relate the mechanics -- diverse network lines, fault-tolerant clusters -- to IT service and business continuity. No matter how beaten up your physical systems are, or how many outages are rippling through the IT systems, the data center continues to deliver mission-critical services.

Michael Kloran

Define parallel processing applications as a way to outrun bottlenecks

WHEN YOU WANT something done quickly, double up the resources working on it.

Parallel processing, also defined as parallel computing, allows multiple server CPUs or processor cores to execute workflows at the same time -- in parallel -- by sharing resources and coordinating actions. Programs are broken down into pieces and then recombined, which completes the work much faster than serial processing.

Parallel processing applications can occur on a single processor with multiple cores, a single server with multiple processors, or across clusters or grids of computers. When set up properly, parallel processing is massively scalable, up to thousands of processors. This is typically seen in supercomputing deployments, not enterprise IT.

While serial processing is good for CPUs, serial communication is better for connections.Peripheral component interconnect express (PCIe) is a lower latency bus than the parallel PCI, and also has higher data transfer rates. Serial means data can travel over the bus in both directions simultaneously, whereas a parallel bus only sends data in one direction at one time. You'll find PCIe connections for network interface cards, graphics cards and storage accelerators on data center servers.

Hardware and applications are becoming more abstracted, with cloud architectures and software-defined data centers. However, a fundamental understanding of how different hardware designs handle different workloads will pay off in the best use of available capacity or a simple hardware fix to a bottleneck problem. For example, rather than rewrite a complex application that is lagging, the data center staff can spec local storage accelerators on PCIe busses. Conversely, when acquiring new hardware for a new application, the data center team can work with developers and programmers to understand how the app will use its resources and plan the best deployment to serve it.

Michael Kloran

Define dedupe as a way to organize, free up storage

BUSINESSES RELY ON a library of customer data to stay competitive, whether your organization sells insurance, IT services or something else. That library gets new incoming data, and changes to existing data, constantly.

With multiple applications using this evolving data from multiple sources, deduplication becomes an essential tool against storage sprawl. Data centers define dedupe as a process to reduce storage demand by saving only one instance of a data point, rather than multiple copies for different systems to access separately. Deduplication lets disparate systems access that single data point.

An example will help define dedupe: Your company keeps a book of data about Customer ABC in its customer relationship management system's database. Customer ABC is also part of a pilot program for your company's new product, which means the same book of data exists in the development group's database for communicating about bugs, new releases and so on. You might also have a separate database used for financials that stores all the same information about Customer ABC. Dedupe means that the company can save just one book on Customer ABC, reducing the storage room required yet not losing any information from which it makes business decisions. The data center's IT systems maintain a reference path so that all three systems know where to find the book on this customer.

There are plenty of practical concerns with deduplication. It can be time consuming and expensive to sort through existing data to match up copies. For new data, storage systems must include the intelligence to synthesize information into one cohesive book; rather than writing a new one each time the company starts tracking a new parameter. But without dedupe, storage demands and complexity would thwart business initiatives like purchasing pattern analytics.

Michael Kloran

Define scalability as the process of growing without changing

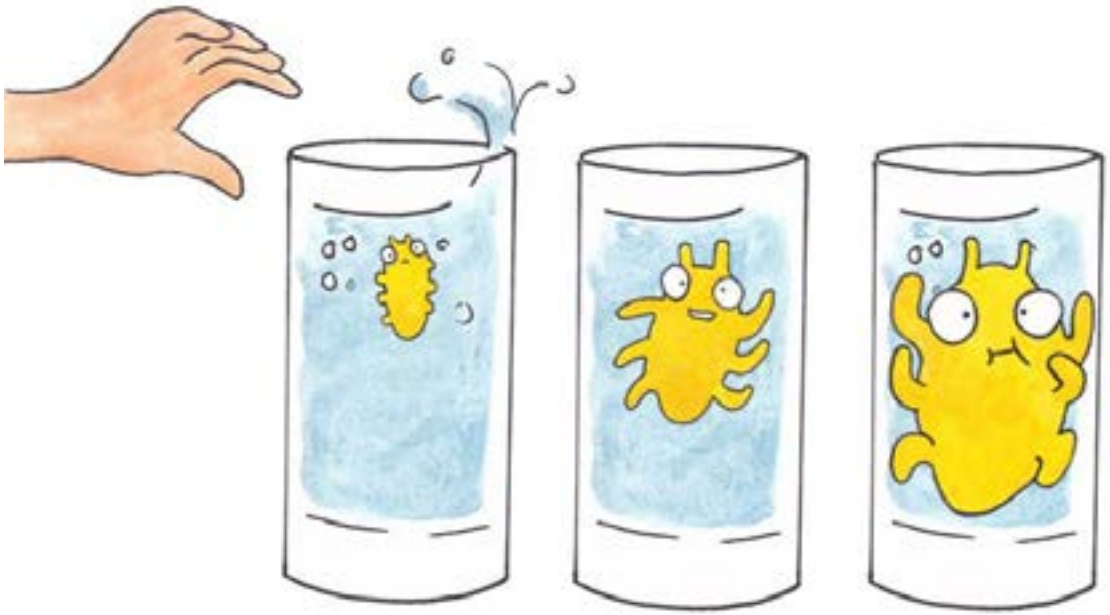

IT WORKLOADS GROW like sponges taking on water: soaking up resources as they become available.

Data center resource scalability gained importance as digital technology started to enable every kind of business. Data center capacity planners define scalability as the back-end IT deployment's ability to grow proportionally with demand, without adding complexity or difficulty to its management.

E-commerce and online transaction processing, as well as content distribution over the Internet, prove the value of scalability. If a T-shirt maker cannot scale up its internal customer relationship management software and its Web server farm to meet increased demand, for example, the business will suffer lost orders and a damaged reputation.

Some IT organizations turn to cloud bursting to handle unexpected or unpredictable short-term rapid scaling up. Data center managers plan capacity to suit the maximum typical demand of IT workloads; beyond that, they rent resources from a cloud provider as needed. The T-shirt maker from our previous example may organically scale up data center servers and storage as its customer base grows, then also burst onto Amazon Web Services' EC2 cloud instances during the busy holiday shopping season. Both methods fit the definition of scalability.

To prove a depth of knowledge in capacity planning, define scalability in contrast with elasticity. Elasticity matches resources to a dynamic IT workload -- when it needs more resources, they scale up; when it needs less, they scale down. This dynamic reallocation is more difficult to achieve with physical servers and systems in a data center than with pay-per-use cloud resources, but with virtualization as well as cloud bursting, it is possible.

Michael Kloran

Define CapEx and OpEx as making sense of dollars

RATHER THAN ARGUING OpEx versus CapEx, it's time to talk about CapEx and OpEx together. Some businesses rely more on CapEx IT spending, others more on OpEx, but virtually all have a mix of both.

In the simplest sense, capital expenditures (CapEx) are single payments in exchange for goods or services, while operational expenditures (OpEx) are ongoing expenses paid to maintain a service.

Business leaders define CapEx as one-time purchases that then depreciate in the organization's care. Data center servers are a classic capital expenditure. Businesses typically buy and deploy a fleet of servers, then depreciate them over a number of years in the data center.

We define OpEx as a recurring expense. Cloud-hosted servers are a classic operating expenditure. The business does not physically purchase and maintain the cloud servers, but rather pays a monthly or per-use fee to access them. There is no depreciation; if a server is not needed, it goes away, and the company may choose to pay for it again in the future.

With OpEx purchases, it's an oversimplification to say that you only pay for what you use when you use it; any administrator dealing with VM sprawl would agree. And without an asset management system, plenty of CapEx spending goes to waste.

Show off your understanding of OpEx versus CapEx as a matter of close IT alignment with business needs. A modern data center employee understands that some IT services are better handled outside of the server racks, while others will be less expensive in the long term with a big upfront investment in hardware and software resources. One such example is big data analytics, where many businesses store and crunch data on owned IT platforms to avoid the costs associated with network traffic to cloud-based big data processing services.

Michael Kloran

Define data center performance metrics to seal the deal

The point of learning about DevOps, change management, IT-business alignment and other hot data center terminology is to improve how back-end technologies and infrastructure perform to meet organizational needs and goals. The proof is in data center performance metrics.

When you're describing familiarity with these concepts at a data center job interview, always present your information in the context of process improvement and best practices. Define some data center performance metrics and how you'd ensure your team met them. Then explain how these performance metrics benefit the company.

For example, don't just point out that deduplication reduces demand for storage. Note that deduplication frees up storage resources for a new big data analysis project that will yield 20% higher customer response rates with minimal capital expenditures.

Research the company where you are interviewing. Are they under industry regulations? Do they value agility and customer acquisition, or do they delve deeply into services for a small group of long-time customers? Approaching the job interview with background information lets you tailor which data center performance metrics you discuss.

Michael Kloran