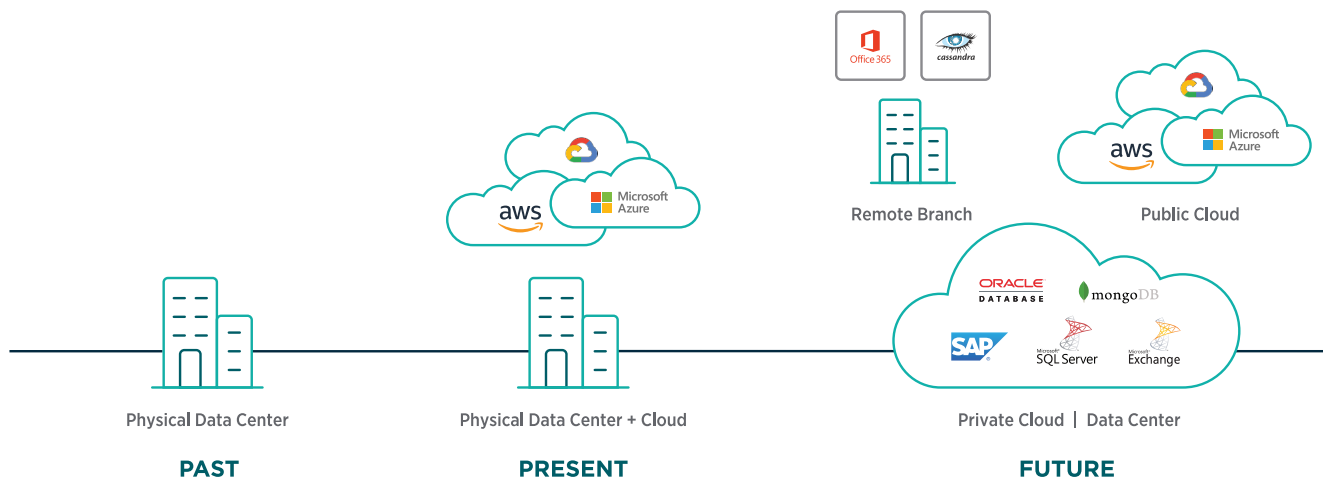


Become Tomorrow's DBA

How to prepare yourself for the new world of hybrid cloud



Your role as a DBA might be changing, but your impact is greater than ever. Let's examine how the new world affects your duties, some of the complexities you'll need to overcome, and the new ways you can add value to your organization.

Everything is moving to the cloud, including databases. As organizations consider cloud or hybrid models, DBAs have an opportunity to re-evaluate responsibilities, practices, and tools to discover new ways maximizing the impact of their role. The benefits of modernization are clear to most enterprises, but the new environment generates additional burdens for DBAs, who face an increase in their workload and require an evolved skillset.

With this advance in technology and responsibility, DBAs expand their value beyond data managers into data stewards, helping their organization innovate the way they interact with, store, and protect information.

PART I: ASSESSING THE CHANGE

As we enter the next era of modernization, organizations and their stakeholders—including DBAs—find themselves thrust into unfamiliar territory. New technologies, services, and philosophies have all sprung forth, creating disruption.

In this drive toward digital transformation, certain areas are seeing dramatic changes. Let's dig into a few of the bigger trends that will affect you.

All Roads Lead to the Cloud

The cloud is one of the biggest identifiable events in modernization. And it's undisputed: databases are moving to the cloud. In June 2019, a [Gartner report](#) concluded that "the cloud is now the default platform for managing data." The report found that 30% of surveyed organizations were storing between 25% and 50% of their structured data in the cloud.

But you knew this already. Cloud adoption and multi-cloud models are the preferred strategies of most organizations. This means a massive shift for DBAs, both in terms of opportunities and challenges.

What can you expect from the new cloud-based approach?



More automation. Fewer tasks to manage as backups, patches, and protection are automated, but the lack of control can be frustrating. Or, possibly, you'll view the removal of menial tasks as an opportunity to focus on innovating the way your organization interacts with data.



Standardization. More can be done with data, within an organization as well as across databases, as structured information becomes standardized in the cloud. As enterprises become more API-driven, standardization allows data to be shared between applications.



Simplification. As you move to the cloud, a lot of the traditional database setup is handled by the platform. You might not get the same level of configuration you're used to, but you've gained the ability to easily add storage and core infrastructure.

Public clouds present the largest disruption in database management. From Gartner's perspective, "on-premises is the past, and only legacy compatibility and special requirements keep you there."

The Impact of Infrastructure as a Service and Database as a Service

The introduction of public clouds has resulted in a boom of Infrastructure as a Service (IaaS) models, including Database as a Service (DBaaS).

With this newfound ability to quickly scale up and scale down services according to their needs, organizations have utilized IaaS to great effect, dramatically increasing their capacity while decreasing the cost of doing business. With lower initial capital requirements, reduced maintenance fees, and fewer IT service costs, organizations have flocked to IaaS as a way to enjoy the benefits of modernization without breaking the bank.

IaaS providers are continuing to work their way up the stack, providing more and more automation for their customers. This push is not only for ease of use, but to ensure best practices are not missed. Included within the IaaS model is DBaaS—on-demand databases hosted on public clouds.

Like other IaaS offerings, DBaaS allows an organization to quickly set up systems for storage and compute in a matter of minutes and free of traditional provisioning errors. DBaaS presents a procurement approach that is only a few mouse clicks away, not the multi-week deployment of an on-premises database.

Organizations are now free to roll out databases according to their needs, with administration and maintenance performed by the service provider. This is a stark contrast to the historical model, in which the DBA has ownership over the entire database procurement and deployment process, along with most administrative and maintenance functions.

But the speed and ease of modern DBaaS could circumvent your involvement, and that has some of your DBA peers cautious.

What does this freedom and flexibility mean for compliance and oversight? More on this in a moment.

Proliferation of Agile Development

Supported by the increased flexibility of self-service databases on the public cloud, agile development is the new trend in productivity.

The fast and affordable database options available in the cloud allow for application development and testing to happen quickly. As previously discussed, spinning up a DBaaS takes only a few clicks, and it can be securely erased at the same velocity. This approach is great for experimentation and innovation because the relatively low cost of a secure environment presents a low-risk opportunity for organizations. Within these cloud-based approaches, developers can take advantage of Continuous Integration (CI) and Continuous Delivery (CD) pipelines. This allows code changes to be delivered more frequently and with more reliability.

This cost-efficient model for production will only further entrench DBaaS and the cloud as a necessary gadget within an organization's toolbox.

The Hybrid, Multi-cloud Approach to Database Management

Public clouds have had a measurable impact on the way organizations manage and utilize data. DBAs like you must now expand your strategies to include scaling multiple databases across the cloud-scape. And, of course, legacy systems are still part of the mix.

In this modern hybrid model, you are expected to be the simultaneous curator of both on-premises and cloud-based databases. It's happening across the industry: 81% of companies have made their DBAs responsible for managing both on-premises and cloud ([DBAs Face New Challenges: Trends in Database Administration, December 2017, DBTA](#)).

As digital transformation accelerates organizations, you should expect hybrid and multi-cloud models as the new norm. Don't be surprised if your domain consists of a series of interconnecting self-service databases interacting with your legacy infrastructure.

PART II: INCREASED COMPLEXITY OF DBA RESPONSIBILITIES

Although methods of collecting, storing, and accessing data have become easier for developers, you are likely seeing the situation through a different lens. Once the master of optimizing data for best overall experience, you must now make sense of a chaotic landscape. Whether it's the sheer number of tasks or the size of data sets, everything has grown bigger and more complex.

The rise of cloud-based models hasn't eliminated existing challenges. Rather, there's a new need to operationalize legacy tools within a modern hybrid model. Added to this complexity, DBaaS and self-service options promote circumstances in which developers can create and access databases without you, often resulting in major operational and compliance challenges for an organization.

The environment has changed, and now everything is more complex.

More of Everything: Data, Databases, Maintenance, and Responsibilities

Today, there's simply more of everything. Naturally, this comes with some good and some bad. Here are a few examples of how everything in the DBA's world has expanded:



More databases under management. Database sprawl is real, and as organizations modernize and expand, so do the number of databases being maintained. But this doesn't mean more DBAs, mind you. As data grows, you could move from managing 50 databases up to 500. In a recent survey, 66% of companies saw an increase in the number of database instances being managed per DBA ([DBAs Face New Challenges: Trends in Database Administration, December 2017. DBTA.](#))



Data sets are growing in size. There's more and more structured data to manage, with 27% of organizations hosting between 1 TB and 99 TB of data, and another 24% managing 100 TB to 499 TB. Expect that number to grow. ([DBAs Face New Challenges: Trends in Database Administration, December 2017. DBTA.](#)) This increase in volume makes tasks like backup and recovery more difficult.



Expanding portfolio of available databases and database functionality. Fifteen years ago, there were just a few options available. Now it's more than just SQL and Oracle; joining the fray are Cassandra, MongoDB, Redis, YugaByte, Neo4j, and a variety of other databases. Each new database has a specific optimized purpose, and it's your job to understand what they do best.



Increased maintenance across multiple platforms. More patching, more maintenance, and more monitoring, all multiplied by the variety of databases being utilized by your organization. It gets complex. Fast. On-premises maintenance will likely take upward of 70% of your day. This is easier with the cloud, where vendor-handled maintenance can reduce your own

workload to 20%. ([DBAs Face New Challenges: Trends in Database Administration, December 2017. DBTA.](#)) Similarly, DBAs are spending upward of 46% of their day patching databases and adding new features to legacy systems. In contrast, DBAs managing their databases in the cloud spend more than half their day on optimizing performance, not maintenance.

You've probably noticed the trend, but the keyword here is "more." As a DBA, you should expect to deal with more databases, more data, and more complexity.

What to Do with All Those Legacy Systems?

Not everyone has the ability to ditch their old tape and disc backups, and some are making a conscious decision to stay on-premises. As developers become more accustomed to the speed and automation benefits of the cloud, the manual process of managing legacy systems can become a huge pain. How can you meet the needs of developers with the restrictions of your on-premises environment?

Infrastructure as Code (IaC) software—like Chef, Puppet, and Ansible—allows you to mimic some modern functionality, replicating the self-service procurement capabilities of the cloud. This gives developers a taste of the speed and agility they expect from the cloud from an on-premises environment. Beyond that, a lot of processes within the legacy environment are custom or manual.

How can you improve operations of your on-premises environment? Currently, automation is handled by custom scripts, but many DBAs find that to be daunting. Not to mention custom scripting can take a considerable amount of time. There is room to innovate to move from this very slow and custom environment to a similar experience to what you would find in the cloud.

Whatever you come up with, your management model needs to incorporate existing systems working alongside the cloud, preferably with the least amount of moving parts.

Overcoming Shadow IT

The last thing any DBA wants to hear is: "Developers are going to continue to choose the database model that is most convenient for them" ([How DBAs Can Survive and Thrive in a World of Agile Development, March 2019. DBTA.](#)) In creating self-service databases, including DBaaS and IaC, developers are circumventing the controls you put in place. This creates a multitude of management issues, for you and your organization, including questions like these:

- **How are databases being backed up?** And while we're on the subject, where?

- **Is the data protected?** Please say that your data is being backed up in case of corruption or disaster.
- **Where is my PCI/PII?** Is your HIPAA data floating out there for the world to see?

In all seriousness, these concerns need to be addressed within every organization, and regardless of who procures the database or where it resides, it falls on you to answer these uncertainties.

From a production-readiness perspective, you are still on the hook to make sure that everything is consistently backed up and recovery tested. Maybe you didn't choose the platform, but you need to be capable of managing any database procured by your organization.

PART III: THE VALUE DBAS PROVIDE IN THE NEW WORLD

DBAs aren't going away. Every organization still needs entities to ensure information is being effectively and safely managed. The question should not be "What is my purpose in the new world," but "How can I build a skillset that extends beyond on-premises, across multiple clouds, and beyond? How can I add value to the data I protect?"

Innovative Ops for the Cloud

As a DBA, you are still required to optimize systems, but now you have the tools and resources to rethink that process. In your new role, you are the pathfinder for a new frontier. Here are a few ideas to spur your imagination:



Data and database guru. Responsible for knowing what exists on-premises and how to best migrate these systems, you are the ultimate resource for selecting the correct service based on the needs of the enterprise. The worst possible scenario is an "accidental" hybrid and multi-cloud approach, in which different groups within the organization act independently in selecting cloud platforms.



Operations innovator. By understanding the proprietary features of certain public clouds, you will be able to avoid future legacy-like circumstances. You should be able to offer guidance on cloud platforms and the services they offer, helping the enterprise determine how it will take advantage of these features as part of a long-term strategy.



Productivity safekeeper. In this final example, you can utilize a multi-cloud strategy to avoid vendor lock-in and mitigate the impact of outages. If one service goes down, you'll still have another up and running.

It's an unintended benefit of experimenting with multiple platforms.

Other opportunities exist as well, but these examples demonstrate the innovative impact DBAs can have on an organization.

Automation Opportunities On-premises

In evaluating your on-premises systems, look to implement automation tools that replicate what is found on the cloud. With a little bit of automation and custom processing, you'll be able to completely transform your organization.

Here are a few areas that are prime for automation:



Discovery. Automatically discover databases as they are added and removed from the database sprawl, including tablespaces, hosts, and clusters. Benefit from a reduced workload and keep your organization protected as new databases come online.



Protection. Manually detecting databases means that you are vulnerable. Paired with automated discovery, you can back up and protect databases as soon as they're added. Additionally, this should bundle **Retention, Archival, and Replication** tasks into a single job, all made easier with automation.



Recovery. Manually, this takes several steps and requires you to access the backup image and logs in order to identify a point in time for recovery. With automation, you will see an improvement in recovery time.

Automation gives your organization the speed and agility to manage database sprawl. And it's not just protection and recovery; automation can help innovate on your operations.

You're probably all too familiar with the bottleneck that accompanies multiple requests for clone creation. By automating the process for self-service creation, you and your organization will see increased flexibility and operational efficiency. In automating clones for testing and validation purposes, for example, developers and stakeholders can move at their own pace, which lifts a major burden from your plate.

Here are a few examples of how automation modernizes clone creation:



Refreshing test and development environments. Some organizations do this monthly, others weekly, and still some will refresh every night. Automation makes it easy and safe.



Testing multi-cloud disaster recovery strategies.

Disaster recovery testing is time intensive and expensive. By automating bi-directional backups, you can quickly and easily spin up a clone to validate the backup at your different sites. It's much cheaper than the traditional route, and certainly a cheaper way of mitigating a major disaster recovery event.



Compliance auditing. The time-sensitive nature of compliance means that there's always pressure to quickly provide data to stakeholders. With automation, you can easily respond to ad-hoc requests for historical data, delivering data to stakeholders in hours, not weeks.



Tablespace and object-level recoveries. Normally, remounting takes days. Instead of retrieving an entire database, you can use automation to perform tablespace and object-level recoveries in a matter of hours. Why recover an entire database when you need only one portion of it?



Validating software upgrades and patches. Start patching systems with data sets that are as close to live as possible. Validation can now be performed by simply spinning up a clone—using near-live data—testing the patch, and then unmounting the recovered database. The patch can now be safely deployed without the usual operational delays and compatibility headaches. There is even an automated path for DBAs to mount backups without the need for additional hardware or storage space. When you're done, simply unmount the database and go on to bigger tasks.

Evolving from Administrator to Data Steward

In understanding where data resides within an organization and how it is being utilized, you take on the most important role of all: data governance. You must now understand what your organization does with data, not just the mechanics of the database. You are no longer in the background, but responsible for data modeling, data security, and performance monitoring in a hybrid system. By focusing on how data moves from one database to another, understanding how data is being used, and managing the data process across the enterprise, you help to engineer the use of information within your organization.

These practices include:

- Data life-cycle management
- Data tiering
- Data archiving
- Replication

Finally, as a data steward, you will serve as an educator within your organization, teaching others how to navigate the new cloud-based environments, demonstrating novel approaches to the data, and promoting innovation throughout the enterprise.

CONCLUSION

Transitioning into a hybrid or multi-cloud management model doesn't need to be a nightmare. By using modern tools and strategies, you can now adopt a bigger role, optimizing and defining the way an enterprise interacts with its data. Rethink the way you've been managing your databases. Understand how data is being used by your organization and offer strategies that promote agility and speed. By altering your approach, you are no longer a database mechanic, but a data engineer.



Global HQ

1001 Page Mill Rd., Building 2
Palo Alto, CA 94304
United States

1-844-4RUBRIK
inquiries@rubrik.com
www.rubrik.com

Rubrik, the Multi-Cloud Data Control™ Company, enables enterprises to maximize value from data that is increasingly fragmented across data centers and clouds. Rubrik delivers a single, policy-driven platform for data recovery, governance, compliance, and cloud mobility. For more information, visit www.rubrik.com and follow @rubrikInc on Twitter. © 2020 Rubrik. Rubrik is a registered trademark of Rubrik, Inc. Other marks may be trademarks of their respective owners.

20200112_v1