



Data Driven Government

Why it's important to take a
data-centric approach to the cloud

Data Driven Government

- 01 | Building Data-Driven Government
- 02 | What Does Dev/Test Really Mean?
- 03 | Personalizing Your Cloud Storage
- 04 | Which Database Service is Right for You?
- 05 | Moving Beyond Backup Tapes



Building Data-Driven Government

Why it's important for agencies
to take a data-centric approach
to the cloud



The amount of data available today is enormous, bringing with it unprecedented opportunities for analytics and innovation. But with such power, comes challenges: Existing legacy systems often have trouble consuming, analyzing and storing the vast quantities of available structured and unstructured data.

Turning to cloud computing—the on-demand acquisition of IT services through the use of shared servers—has been the logical next step for forward-looking agencies. The versatility, scalability, elasticity and reliability of the cloud allow agencies to reduce cost expenditures while simultaneously improving efficiency and analytics capabilities.

“Spinning up compute resources in the cloud has a direct analogy to the physical clouds,” says Kirk Kern, U.S. Public Sector Cloud Chief Technology Officer at NetApp. “The clouds amorhously change. They’re always different—they’re there one day and gone the next—and that has some similarities to the way services in the IT clouds operate. Applications must be able to start, stop, restart or scale their resource consumption to accommodate the dynamic nature of clouds.”

Because cloud computing relies on the use of shared resources, agencies using the cloud can dial up or down services to meet their computing needs. In times of high usage—such as when testing a new application or releasing a new service—agencies can scale up their computing power, and then scale it back down when the demand for computing power decreases. This means that agencies no longer have to invest upfront in the physical hardware to meet their peak computing needs, and then maintain these systems even when they are not being used. Instead, computing power can match demand, providing flexibility and agility, as well as enormous cost savings.

As with any new tech solution, organizations often wonder about the security of their information. This is especially true with government, where a network breach could compromise anything from national security to the privacy of individuals.

“Each organization needs different data protection and has different organizational and risk tolerances,” Kern says. “As a result, the agencies have to perform an analysis before adopting these cloud technologies, especially if it’s going to be a public cloud or even a hybrid cloud solution.”

To best use cloud while ensuring that sensitive information is kept safe, companies like NetApp have developed a cloud data fabric—a technology that brings together computing systems in order to give organizations a unified way to control their data, providing them data stewardship. By stitching together the resources of multiple clouds, agencies can leverage their elastic and dynamic computing resources while still maintaining complete control over where and how data is moved.

“The government can choose between on-premise and off-premise IT assets, or even some combination of both commonly referred to as hybrid, to achieve operational activities like elasticity, scalability and disaster recovery,” Kern says. “Fabrics offer even better protection.”

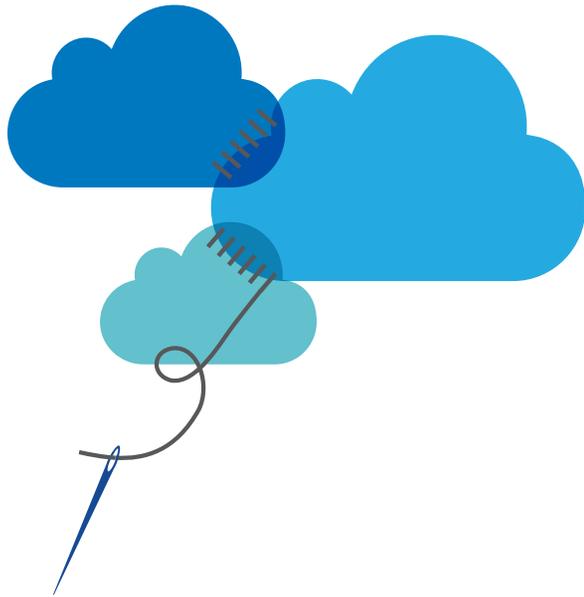
This level of control allows organizations to ensure sensitive data is kept safe. It also allows experts to monitor who has access to what information, protecting data by preventing unauthorized users from ever being able to access it in the first place.

“What we’re starting to see is techniques for Data Governance being employed to create metadata, or metadata repositories to track data locations, data ownership and even the migration paths. A simple question would be: Can you tell me where this file is; who edited it last; and what’s the most recent version? Can you show me all versions of this that exist in multiple clouds? And so that’s the type of capabilities, data protection and data stewardship that a data fabric can provide,” Kern says.

Despite the benefits of cloud, adoption can seem daunting. Properly implementing a cloud solution requires careful planning and strategy even before any service is implemented. A good cloud strategy has a comprehensive approach for placing data into the cloud and moving it back out again. Agencies need to determine which security mechanisms they need and understand that maintaining data stewardship is critical to supporting the dynamic nature of service consumption.

Perhaps most critically, agencies need to think about where they can derive the most benefits from cloud.

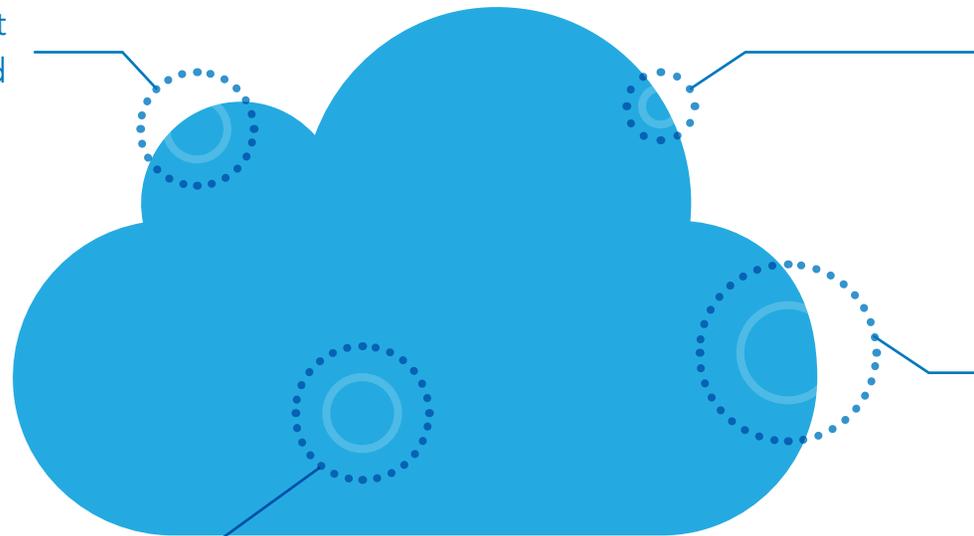
“Organizations tend to pick off the low-hanging fruit, but the thing that cloud really brings to bear is the ability to command vast amounts of resources and dial them up quickly,” Kern says. “Rather than looking at simple workloads, go after workloads that will give you dramatic returns on the investment you’re making in the service.”



By stitching together the resources of multiple clouds, agencies can leverage their elastic and dynamic computing resources while still maintaining complete control over where and how data is moved.

Putting test and development workloads in the cloud

Identifying which database service is right for you



Modernizing your backup and recovery plans

Leveraging private storage for public clouds

Throughout this e-Book, we will examine how to build out cloud solutions for government. Specifically, we'll focus on four areas:

- › Putting test and development workloads in the cloud
- › Leveraging private storage for public clouds
- › Identifying which database service is right for you
- › Modernizing your backup and recovery plans

Through agency examples and interviews with IT professionals, we'll show you both what government is doing today, and what it is capable of looking forward.

What Does Dev/Test Really Mean?

Move workloads to the cloud
and take advantage of the true
meaning of dev/test



Developing and testing new applications is important for any successful organization, especially in the public sector. In order for government to continually offer the best, most efficient services possible, agencies need to be able to innovate at-or-above the pace of private enterprises.

Dev/test is one of the most common workload types being moved to the cloud, but still many struggle to understand just what exactly this means from a practical perspective. In fact, it's one of the most common questions that Kevin Hill, cloud solutions architect at NetApp, finds himself answering.

"The greatest benefit of cloud is agility and flexibility," Hill says.

While these are boons across the board, they are especially important for dev/test workloads. Using the cloud, agencies can quickly spin up their resources, test out new applications and react quickly to both successes and failures.

"Dev and test is a really good workload for cloud usage for a couple of reasons," says Mark Ryland, chief solutions architect at Amazon Web Services. "First of all, the speed of development is an immediate benefit for a development team. You don't spend any time acquiring equipment or software—the things you need to build systems. Sometimes it could take months just to get the stuff you need to start developing software, and that can be reduced to days, at most weeks, using cloud. Right away, the development process gets off to a much faster start."

Beyond that, agencies doing dev/test in the cloud are not limited by the hardware that they have. This is especially important because, when developing, it's often difficult to know from the get-go just what systems, hardware or languages will work best. Cloud eliminates this problem by allowing experts to experiment and try different things.

“You don’t have to make a big decision in advance—like on using java or python. You can actually try both, have parallel teams working, and then decide what the best language and frameworks are,” Ryland says.

There is also a significant cost benefit associated with using the cloud for dev/test workloads.

“Often development is not a 24/7, 365 process,” Ryland says. “You come in Monday morning, work for a while, and shut your stuff down at the end of the day. That’s highly efficient in a cloud environment because you’re not paying for the facilities that you’re not using on the weekends and at night—you can shut it down and save a lot of money. In the development phase, that kind of elasticity is extremely beneficial.”

When it comes to the testing phase, the cloud can help agencies both realize savings and test their newly developed application in the most realistic environment possible.

“You can actually test at scale—test realistically, with production-level loads, testing that is really difficult in an on-premises environment,” Ryland says. “In an on-premises world, to test with realistic loads, you’d have to buy equipment equivalent to what your production system works on and have it sitting idle for most of the time. That’s just not realistic. You don’t have to do that in a cloud environment. You can literally stand up a copy of your production system, run it for a few hours or days, do a full-scale load testing and then shut it down. It’s a completely different way of solving a bunch of problems that have always plagued people when they’re doing software development and test.”

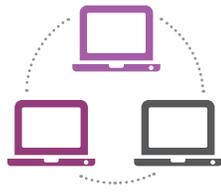
The very premise of dev/test involves an element of experimentation, and by greatly reducing the costs associated with project implementation, agencies can turn failures into an informative, necessary part of innovation.

“If it’s unsuccessful, agencies can say ‘nice try,’ document what happened, and then with the cloud basically release the resources. There’s no longer any chance that you’ll have idle resources or assets just sitting there waiting for the next project. There’s no longer any charge or penalty for that particular failure,” Hill says.

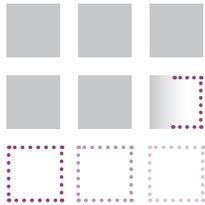
When the project is successful, however, the rewards are even greater than they would be before cloud implementation. For example, the flexibility and elasticity of the cloud allow agencies to test algorithms at large scale—something that they would be hard-pressed to do in a traditional dev/test environment.

```
0100001011
0001100101011011
00111100010100000001
01001100101011011100010
000001110011011010110
110000101100001000000110
100011100100111100111100
001000101011101001100
1000010010000011
000110110000
```

Data



Computing



Automation

“If it’s successful, there’s a lot more capabilities, so you’re rewarded a little more,” Hill says. “You can either keep that project in the cloud and take advantage of different cloud assets—such as being able to flexibly spin up as many resources as you need or, if needed for policy or other reasons, move that application stack with its compute and its data into a private datacenter.”

From a practical perspective, Hill breaks down dev/test into three key areas of consideration: data, computing and automation. Agencies need to ask themselves what they are testing and if they already have the data to leverage it. The goal is to understand what data is available, and how best to work with this information to achieve the desired result.

Next, agencies need to address the compute portion. By choosing a public cloud, agencies can rapidly and efficiently spin up, spin down or resize their computing power to meet the individual dev/test requirements of a particular organization or project. Finally, agencies need to consider a level of automation and orchestration that can make the implementation of these services easier.

Services such as NetApp Cloud ONTAP can help with this by providing multiple storage consumption models and an easy interface through which to control public cloud storage resources. And, from a timing perspective, it’s infinitely quicker to stand up Cloud ONTAP systems than physical ones. This can give agencies the agility to experiment with new dev/test applications without a lag time between ideation and creation.

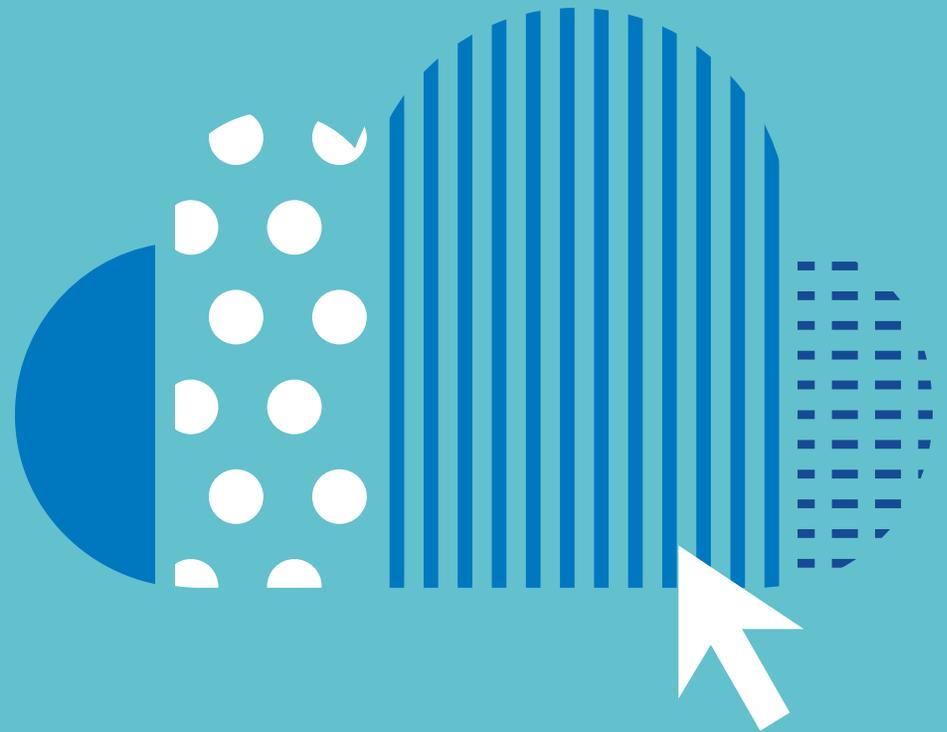
“Even a small, entry-level physical system would take weeks to stand up, but in roughly 20 minutes you can have a fully functioning, up-to-date, Clustered Data ONTAP system running in your Amazon account,” Hill says. “Within one day, theoretically, I can actually run 10 tests off the same type of configuration after I replicate my data up to the cloud, and then just as easily replicate that data back out of the system if I want to keep it. That agility and flexibility is just massive.”

Cloud expands the possibilities for dev/test by providing a scalable, flexible, agile and inexpensive environment that facilitates innovation, allowing agencies to optimize application development and produce the best products at minimum cost.

From a practical perspective, Hill breaks down dev/test into three key areas of consideration: data, computing and automation.

Personalizing Your Cloud Storage

The value of having a flexible,
personalized storage solution



Although agencies across government are migrating to cloud, nearly 89 percent of IT professionals feel some level of apprehension about this process, **according to a recent study done by Meritalk.**

Much of this unease stems from questions of data security and flexibility. Government leaders want to realize the flexibility, scalability and increased collaboration of the cloud, but wonder how they will be able to maintain complete control over their data if they release it entirely to the public cloud.

“They’re concerned about security. They’re concerned about how flexible any public cloud is going to be once they start using it,” says Demas Abraha, CTO of Blue River IT. “How can we have our secure end-to-end encrypted data, and optionally ensure that access to physical storage is within reach of the end-user and the owner of the data? How do we do that?”

In response to these concerns, companies like Blue River IT have developed private storage solutions for public clouds. But, what exactly does this mean?

“Whether you’re talking about performance of your services security, flexibility, or even managing the cost of offering your storage, you have to start from the data,” Abraha says.

Essentially, private storage for public cloud allows agencies to harness the elasticity and cost savings of using a public cloud, while retaining the performance, availability and security associated with using dedicated enterprise storage. It does this by connecting on-premises storage to the public cloud at a colocation facility, thereby allowing organizations to place their data “next to” versus “into” multiple clouds.

“You can take advantage of cloud compute while retaining ownership of your data provided by the private data storage,” Abraha says. “In order to continue existing data security practices because you own your data, you use private storage and take all the advantages of using public cloud such as utility computing.”

For agencies with sensitive workloads, this eliminates anxiety around security and privacy while simultaneously improving their computing power. Agencies can employ end-to-end encryption, ensuring that their data is safe both when in transit and at rest. And, by using these solutions agencies can retain the ability to have complete control over who has access to what type of data.

“Think of privacy issues,” Abraha says. “If you have healthcare-related data, how do you ensure that only people who need to have access have access to it? So yes, it’s all under the umbrella of security, but the privacy and the flexibility to transfer custody of data from one entity to another is what I’m talking about.”

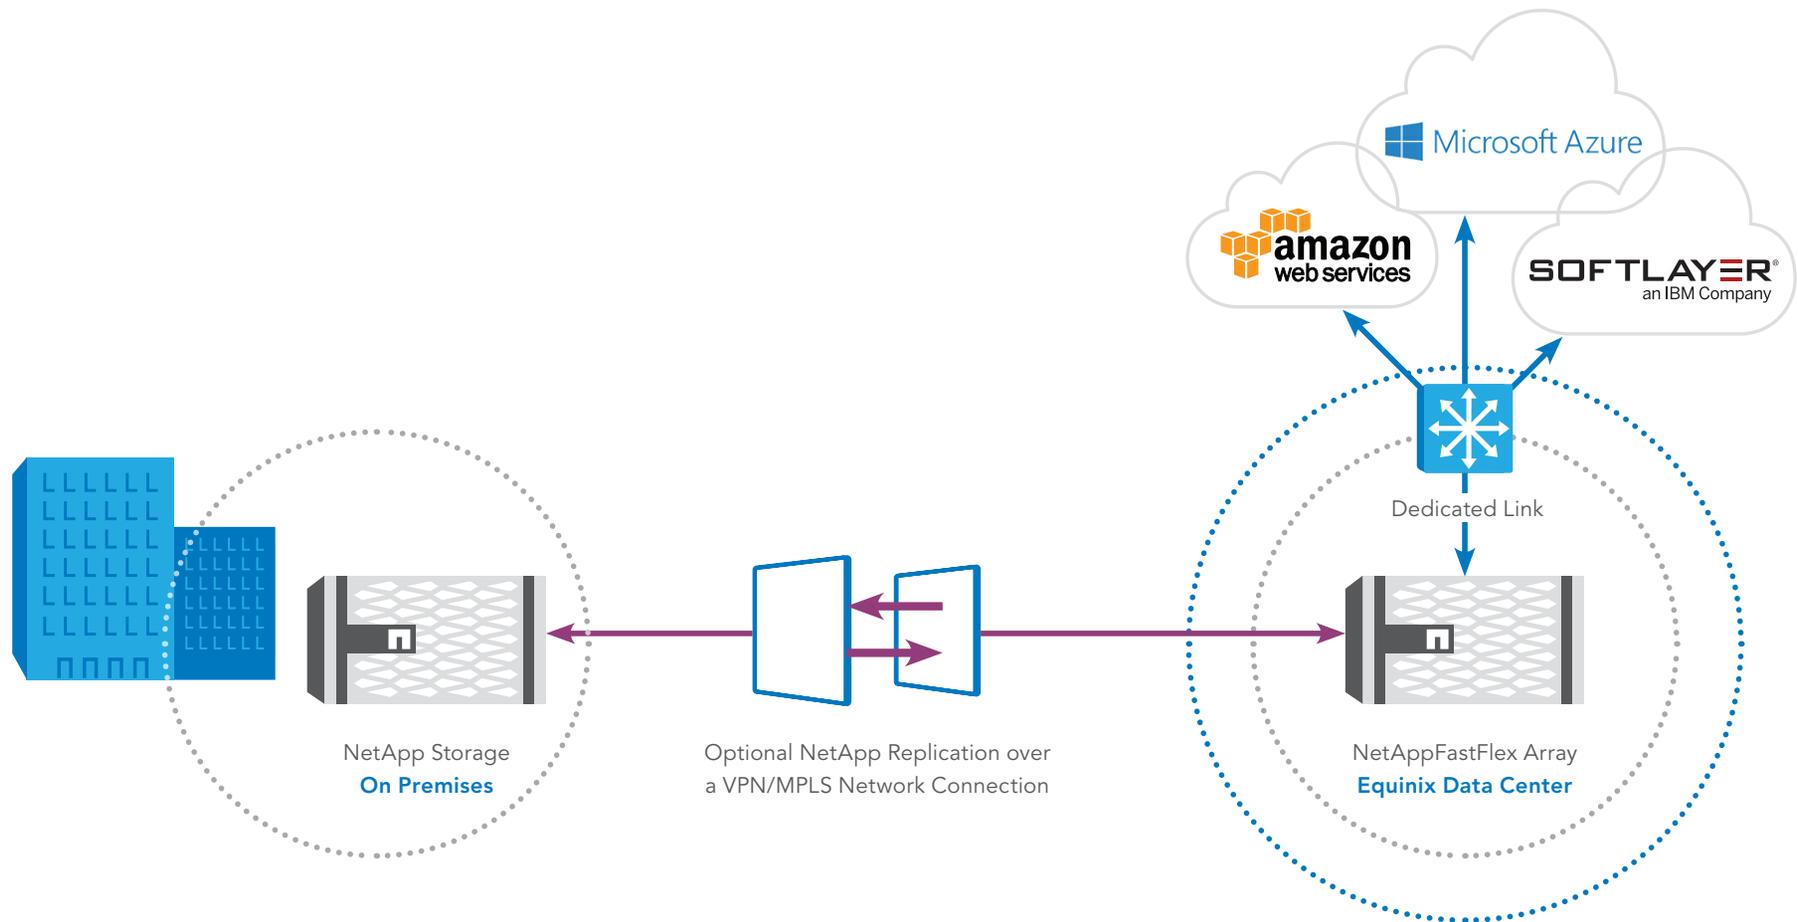
Blue River IT’s Mallard Solution is one example of a solution that capitalizes on these capabilities by delivering an integrated on-premise (or co-located) cloud platform, where data management tools can be used to manage cloud-based data and various environments can be leveraged to best meet an organization’s requirements. Because there is no way to circumvent latency associated with physical distance between the data centers where private storage and public clouds are located, Blue River IT looked for colocation facilities in close proximity where the latency would be 4 milliseconds or less.

“Doing that, we essentially made NetApp storage to be part of Amazon’s compute,” Abraha says. “You can continue to use the tools we spoke of earlier as if they are in native storage in the public cloud. So you have the flexibility of the private cloud, with the public cloud advantage of compute.”

By facilitating strong data governance practices—an element that should be key to any agency’s IT planning—agencies can know where their data is, how it’s moving, who has access to it and what measures will be taken in case of emergencies or unexpected situations.

“Whenever we talk about data, all these multidimensional requirements always overlap,” Abraha says. “So whether it’s security or privacy, all of them have to do with data governance... Once you get to the cloud, if you make a mistake in your strategy, it can be very hard to fix it without major expenses and difficulties. Whereas if you start with data governance, you’ll know things like: ‘What happens to my data if I were to do this? If these things were to happen? If the cost were to go beyond what we expected?’ This is one feature of data governance: knowing what happens and how we can minimize risk.”

An equally important, but perhaps less talked about advantage of using private storage options for public clouds is that it opens up agencies to a world of new technology possibilities allowing for workloads that might fall outside of the scope of a particular cloud provider.



With NetApp Private Storage, once you make your first cloud connection in an Equinix colocation facility, it's easy to add new clouds or switch clouds without moving your data.

"If you're using public cloud for your storage, then basically the new tools, ways and methodologies have to be supported by those cloud providers, whereas private storage gives you the possibility of employing newer technologies as well," Abraha said. "So it enables you; it gives you the flexibility to acquire newer technologies, newer solutions, and multiple cloud providers."

What this means is that private storage for public clouds encourages unprecedented flexibility and innovation at the same time as it delivers top-notch security. Agencies no longer have to be concerned that, when it comes to cloud, they have to sacrifice either security or capabilities. Instead, they can focus on using the technology to meet their mission objectives.

Which Database Service is Right for You?

Effective ways to save time and money with Database as a Service (DBaaS)



It's a common set of questions many developers and IT managers face:

- › Which database technology is right for my application and my agency's mission?
- › Do I need multiple database engines?
- › How do I acquire and afford the necessary expertise to manage multiple database platforms?
- › How do I avoid technology lock in?

These can be tough questions to answer, but powerful new options have emerged.

Database as a Service (DBaaS) is a data management platform comprised of automated software tools, industry best practices, application blueprints and supporting technology services that are offered by cloud operators, in public, private and community deployment models.

The DBaaS platform supports complex data intensive applications without the need to hire a database administrator to support the end users or application teams. DBaaS allows IT professionals to create solutions much more quickly and with far fewer defects, while supporting a variety of database solutions based on SQL and NoSQL database engines.

Federal, state and local governments are always seeking to save time and money in order to deliver services to their constituents. In doing so, they must first consider their mission and operational goals, and prioritize their budgets to meet those goals. They are, therefore, always in search of IT efficiency, business agility, high application availability and simplified data management opportunities.

One of the first steps towards achieving these larger goals is to utilize a cloud platform to drive down infrastructure costs. This provides a huge return on investment through reduced server equipment and administration labor costs. The next logical step is to employ DBaaS to drive down labor further, but more importantly by removing massive bottlenecks to application deployment and business agility.

Another area for consideration when looking to a DBaaS solution is storage. To avoid a business impact, your database storage must deliver sustained performance for multiple instances-

whether on premise or in the cloud. By utilizing technologies such as **FlexClone** and **Virtual Storage Tier** from NetApp, DBaaS providers can bring increased customer functionality and decrease the administration overhead of managing databases.

“The Database as a Service allows managers and administrators to decide how and why a new application server should be deployed without being limited by the availability of a skilled or at least competent DBA,” says Sean Jennings, senior vice president of solutions architecture for Virtustream.

“It used to take weeks, if not months, to get a new server up and running,” he says. “But, with Database as a Service, deployment times for new cloud-based systems can be reduced to hours or just minutes. It could take weeks or months to get a new server in the old days. Now a server can be deployed in moments and ready for the application team in minutes.”

DBaaS gives you the ability to rapidly and reliably deploy new databases that adhere to established best practices while eliminating standard wait times. If you’re a developer, this capability gives you a self-service option for launching your next application. The “drudge work” of the database administrator is no more, Jennings says.

“The database administrator used to be the bottleneck to deployment,” he says. “Now, a developer can use a simple dashboard that walks them through several key questions which then launches an automated process resulting in a production ready database solution. This is a huge force multiplier.”

Even an end user can use the dashboard to walk through a very simple provisioning process that is point-and-click, Jennings says. “You don’t need to be an expert in Oracle SQL. You follow a few questions and it will guide you to a complete solution that meets all of your requirements from performance to security and compliance.”

This is particularly useful for government, which needs to meet certain rule sets or standards that guide database deployments. Many organizations, including agencies of the federal government, require compliance-driven solutions that often face personnel bottlenecks. Using the dashboard approach, application developers can consistently and reliably set-up a database, taking into consideration everything from how storage is laid out and managed, to role-based access controls for security purposes. The DBaaS engine then ensures that the established policies for security, data integrity and retention are enforced. The service can therefore save an agency time and money.

“If you’re the CIO or CFO, you’ve now gotten to the point where you’ve taken a lot of the cost out of the equation, in-terms of man hours it takes to deploy,” Jennings says.

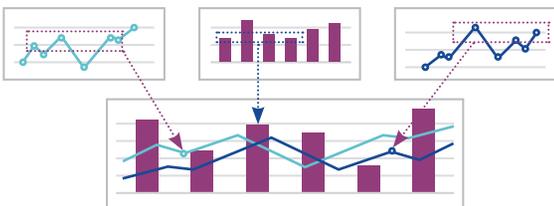
1

Marketing Budget	IN	OUT
Event A		\$15,000
Social media ads		\$1,000
Planning software		\$500
Product launch	\$2,000	

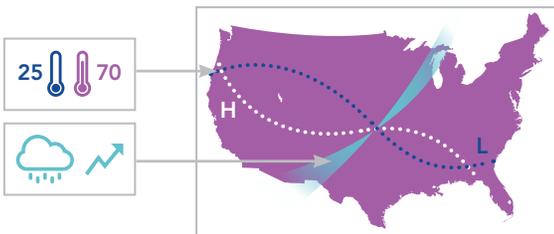
+ \$700

- \$300

2



3



Examples of (1) transactional, (2) analytic, and (3) real-time database systems.

When it comes to the available data solutions, there are three main categories of databases to consider:

- 1. Transactional systems** involve user interaction with an application supported by the database and are usually measured in terms of end user response time. Examples include accounting systems or enterprise resource planning systems that are updated or queried interactively throughout the day .
- 2. Analytic systems** work to solve “big picture” questions and detect patterns in disparate data sets. Analytic data systems analyze data from multiple sources and often involve complex and often unpredictable “what if” scenarios in an effort to gain operational insights. What typically happens is that many different views into the data are created. The reports from these systems are typically batch data oriented, meaning they use multiple data sets, and can require several days to complete. Many of the “what if” scenarios are never anticipated by the data engineer or architect, so the processing times can vary greatly depending upon the data size and interrelationships.
- 3. Real-time systems** are an area government is in the early phases of adopting. These systems make use of real-time data inputs or sources, typically stored in memory, in order to rapidly evaluate large amounts of data in real time and analyze emerging trends. An example of this includes monitoring of environmental sensors to forecast severe weather and storms. Another example is the use of real-time monitoring of global financial systems to detect fraud or other nefarious behaviors.

“Real-time has started to move into government as we’ve developed in-memory databases, which allow for much more rapid analysis of data, very similar to what happens with some of these analytics systems,” Jennings says. “So now, you can do the ‘what if’ scenarios where all of the content is in-memory. All of the sudden you can do in five minutes, what used to take several days and even weeks.”

Real-time systems can help intelligence agencies to identify and prevent terror threats and protect critical infrastructure from emerging zero day attacks. In-memory databases can ingest and process massive amounts of information and provide real-time analytics.

One of the more popular in memory database solutions on the market today is SAP HANA. This is a new technology platform that is early in the adoption cycle for government agencies. A wide range of public entities —including technology, defense and health related agencies— are already benefiting from its new capabilities.

SAP HANA is natively a column-based database solution with an in-memory architecture to store data; it includes an abstraction layer to analyze and present data for row-based SQL and object-based queries with traditional tools and skills.

This architecture enables real-time data analysis and, according to Jennings, is already delivering tangible and powerful results for government.

“Government is still dipping their toe in to try it, but the results are there. And, because solutions like SAP HANA are so new, agencies are often finding unexpected solutions and benefits. Therefore it is important for agencies to experiment with these new capabilities and consider their options and potential use cases,” Jennings says.

Government is not always the first to adopt cutting edge technologies, but already there are cases being made for why real-time data is a necessary and powerful tool.

“Think weather systems, transportation, air traffic control, medical and drug research or even tracking military and intelligence threats,” he says. “In each of these cases, government is using real-time data to analyze a set of problems in ways that were previously unavailable.”

Moving Beyond Backup Tapes

Leveraging cloud for efficient and secure data practices



From monitoring defense satellites to examining health information, and financial trends, the federal government has a lot of data. On any given day, data is collected, categorized and analyzed, delivering key insights to agencies and helping agencies turn ideas into reality. But all this information needs to be stored and kept safe in case of a disaster.

Backup and recovery is critical to the successful operations of federal agencies. The loss of valuable data can set an agency back years in terms of research and operations—an unacceptable loss of time and resources.

That's why, across government, agencies are upgrading to SteelStore, a cloud-integrated solution for backup and recovery provided by NetApp.

"Typical backup flows include magnetic tapes, which had to be off-sited and taken to a secure location in case of disasters," says Rich Faris, SteelStore director of product management at NetApp. "The tapes often failed, so you'd bring back a tape to recover and potentially have data loss. And then there's a lot of cost management of tapes. Since they're physical objects, they have to be labeled, boxed up and taken off site."

Put simply, a tape-based recovery method is inefficient and relies heavily on human labor. As the tapes age, data loss is possible—especially if tapes are improperly handled. And, because the tapes have to be physically shipped, there is the risk of the physical media being lost in transit.

"Often people using tapes don't even encrypt them," Faris says. "You'll see in the newspaper a tape falling off a truck, sitting there on the pavement. If that has all your social security numbers or credit card information on it, it hurts all of us if it gets out."

Securely and safely sending backup data encrypted into the cloud eliminates this problem completely. Once in the cloud, the data will remain there until you need to access it, and can be protected with the newest security technologies available. Moreover, data in the cloud can be stored across regions, so the loss of any individual data center does not mean the loss of data. This is especially important when it comes to disaster prevention.

“If the primary site or any site is rendered inoperative or destroyed, the data is still in a different location, and thus even though they lost the site, they can bring the data back and not lose it,” Faris says. “The cloud is often also replicated in multiple places, so data gets put not just once in the cloud but is also replicated in multiple centers, such that even if one of the data centers of the cloud is destroyed, the data will continue to be available and secure.”

Regardless of the motivators behind backup and recovery, a cloud-based backup and storage system is also much faster than traditional solutions.

“With tapes, you’re getting your data back at the speed of a truck. We’re getting our data back at the speed of the internet, and are able to reach out and grab the data right away without waiting for the truck to bring the tape back,” Faris says.

Cloud-based backup and recovery systems get rid of highly labor-prone activities and speed up data recovery. Solutions like SteelStore can also de-duplicate data in order to save network bandwidth, space in the cloud, and associated cloud billings.

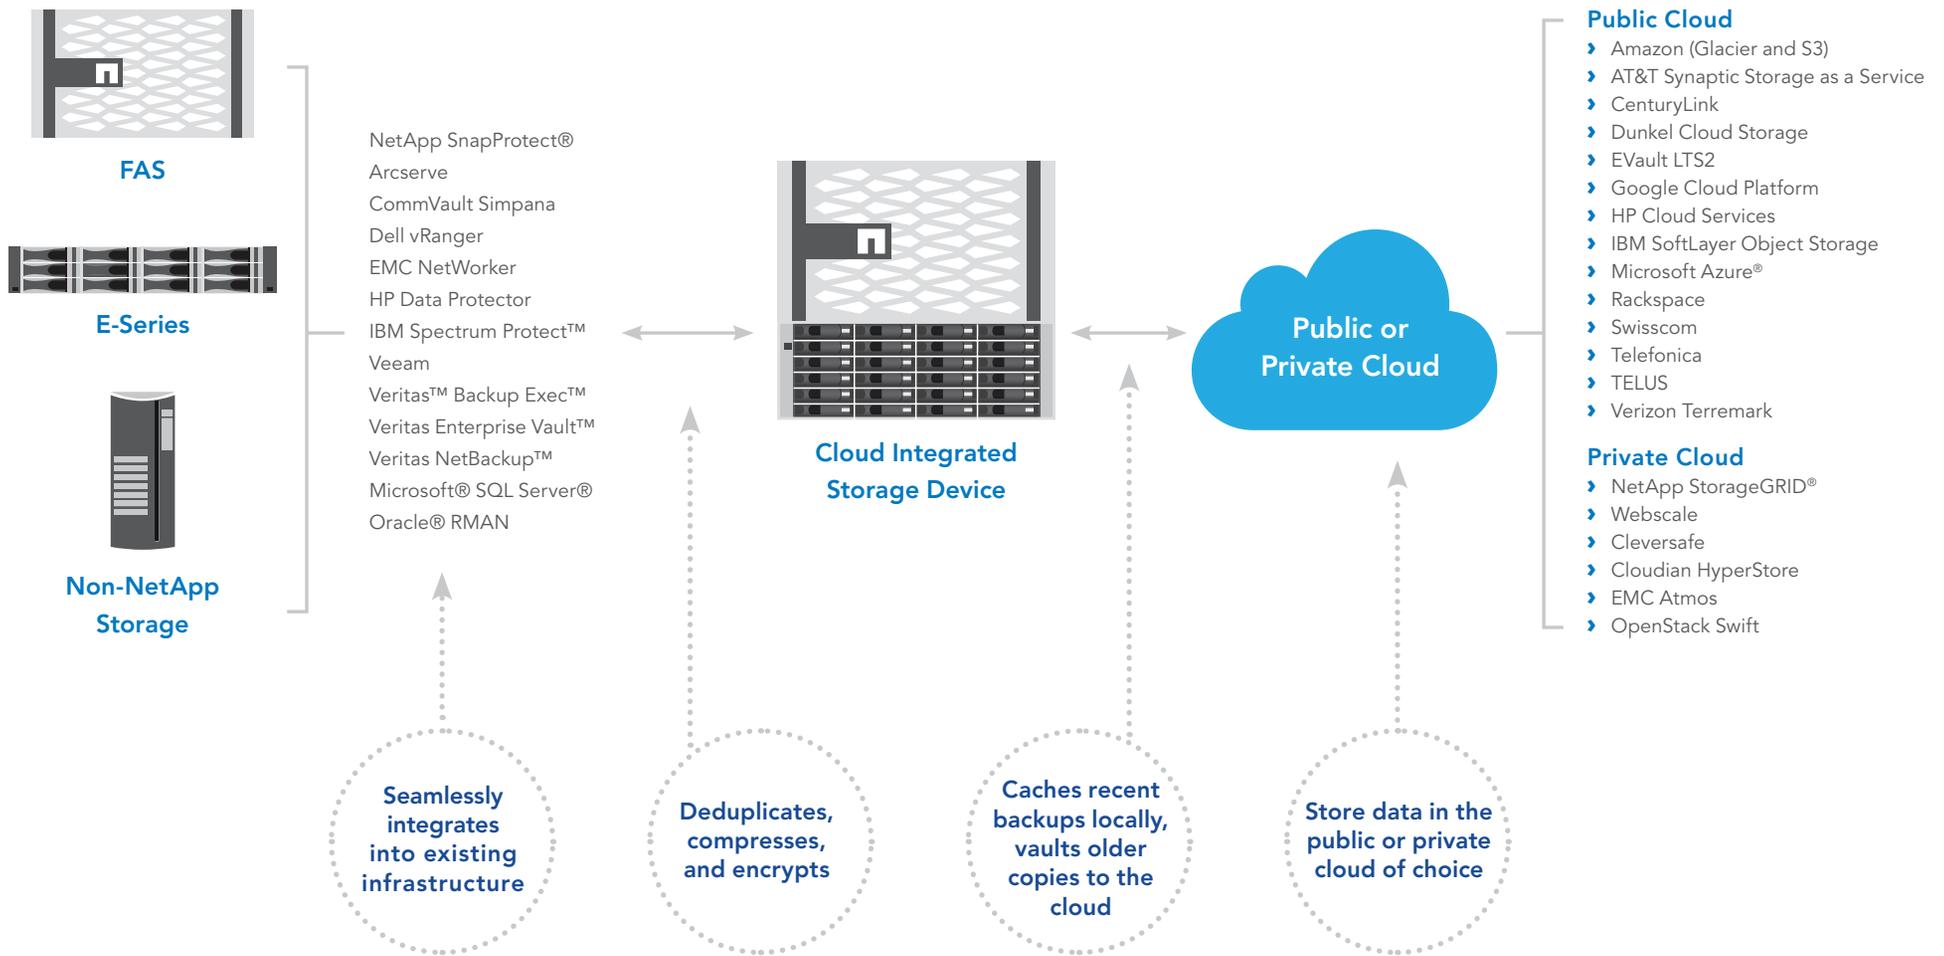
Perhaps most critically, using SteelStore technology allows agencies to realize the benefits of cloud without rendering void their investments in previous backup and recovery systems.

“The other big value of the SteelStore device is that we don’t replace existing infrastructure,” Faris says. “It’s not a rip-and-replace device, so instead of making the customer remove their entire backup infrastructure and replace it with us, we’re actually sitting in as a drop-in replacement. This means that they can preserve their investment in those tools and all that effort putting together backup scripts, and still can make use of the cloud and efficiency thereof.”

This solution has proved highly successful across the public sector, helping improve the backup and recovery process for federal, state and local government organizations.



The cloud is often also replicated in multiple places, so data gets put not just once in the cloud but is also replicated in multiple centers, such that even if one of the data centers of the cloud is destroyed, the data will continue to be available and secure



SteelStore provides seamless integration with leading applications and cloud service providers

“This type of technology, combined with the utility of cloud storage, will reduce costs by allowing us to eliminate tape backups and associated operational complexity, while giving a more robust backup solution,” said Reid Swick, system architect in King County, WA. King County uses the cloud for storage and compute, and in the past year has started using SteelStore to backup to the cloud.

Building out a backup and recovery plan, or updating an existing one, is a lot of work—especially for government, which handles enormous quantities of highly sensitive information. Unlike other workloads, backup and recovery is more of an insurance method in that it does not immediately add value for citizens the way that new products developed through advanced dev/test methods might. But that doesn't mean it's any less important, Faris cautions. In fact, proper practices can actually free up experts to work on new innovations instead of simply maintaining existing, inefficient systems.

"Everybody buys insurance because they must, but it's not a function that actually adds value to their end-customers as much as some other things that they could work on," he says. "But there's a lot of cost of management of tapes, and new technology can actually allow organizations to turn their administrators and have them work on higher value products for their customer base."

Across the public sector, organizations have recognized the benefits of cloud-based backup and recovery, and are undertaking this effort. Many are far along in this process. Before spending valuable time and resources investing in new technology, Faris recommends reaching out to other organizations for tips and advice.

"Talk to other agencies," says Faris. "You might be surprised with the level of adoption that is happening out there, and you can leverage some of their best practices. Because these other agencies, they put a lot of effort into it because they see a lot of promise. So take a look at what they're doing, talk to them, and then make your own decisions based on what you hear."



Learn more at www.NetApp.com