A SOLIDFIRE PAPER

Designing the Next Generation Data Center

NetApp®

Table of Contents

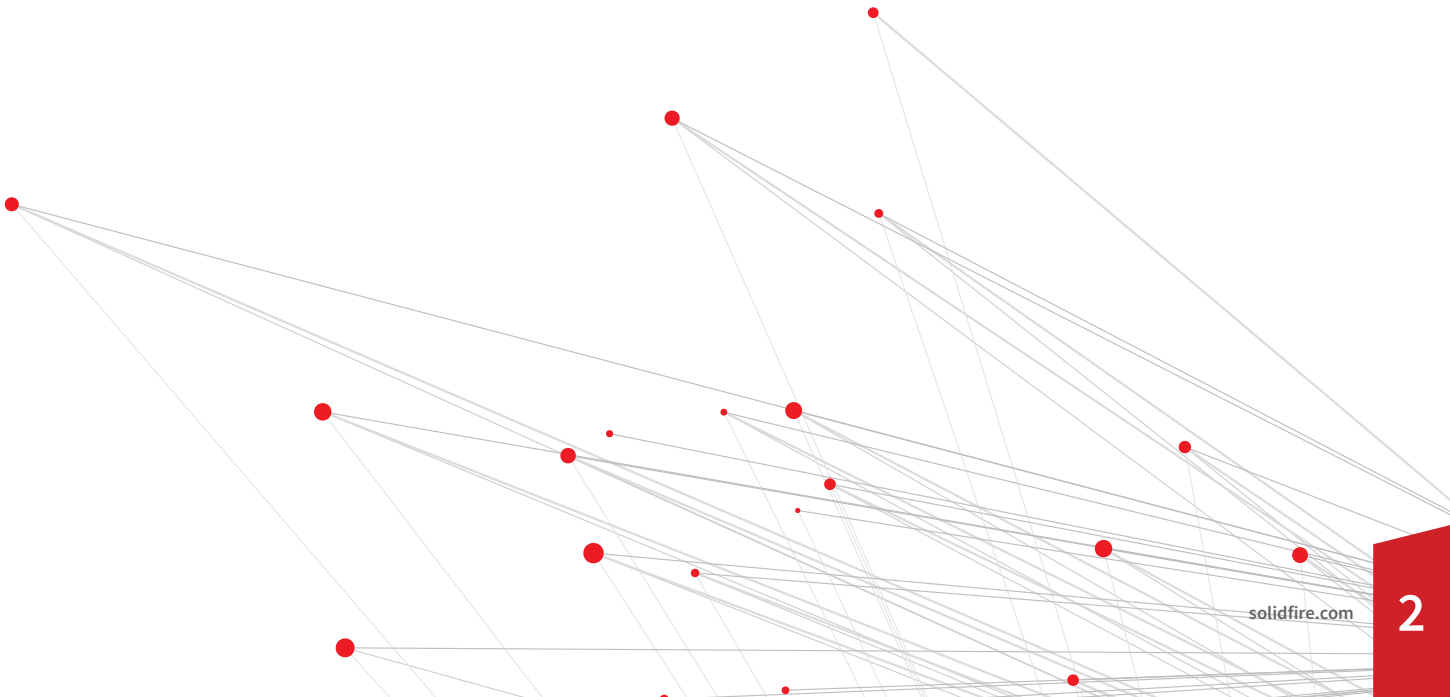

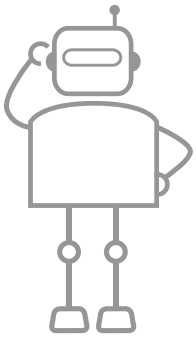

Why Read This Guide?

At the turn of the 21st century, virtualization promised a more flexible, centralized data center. But the global IT community wasn't aware of that need. Early adopters realized the potential business advantages well before it gained mainstream adoption. Today, virtualization is pervasive, reaching market maturity in its lifecycle. Cloud computing, however, is on a growth trajectory. Consider the following studies:

- Goldman Sachs forecasts cloud computing and infrastructure spending to grow at a 30% compound annual growth rate from 2013 through 2018, compared to a 5% growth for enterprise IT overall.[1]
- IDC forecasts public cloud spending will more than double to $127.5 billion by 2018 — $82.7 billion in SaaS, $24 billion for IaaS, and $20 billion in PaaS.[2]
- Computerworld's 2015 forecast study revealed that 42% of IT decision makers planned to increase spending on cloud computing in 2015.[3]
- Cisco's global cloud index report predicts that by 2018, more than 78% of workloads will be processed by cloud data centers; 22% will be processed by traditional data centers.[4]

1. Forbes. "Roundup of Cloud Computing Forecasts and Market Estimates,2015." Available at http://www.forbes.com/sites/louiscolumbus/2015/01/24/roundup-of-cloud-computing-forecasts-and-market-estimates-2015/.
2. IDC. "IDC Forecasts Public IT Cloud Services Spending Will Reach $127 billion in 2018 as the Market Enters a Critical Innovation Stage." Available at http://www.idc.com/getdoc.jsp?containerId=prUS25219014.
3. Forbes. "Computerworld's 2015 Forecast Predicts Security, Cloud Computing And Analytics Will Lead IT Spending." Available at http://www.forbes.com/sites/louiscolumbus/2014/11/26/computerworlds-2015-forecast-predicts-security-cloud-computing-and-analytics-will-lead-it-spending/
4. Cisco. "Cisco Global Cloud Index: Forecast and Methodology, 2014–2019 White Paper." Available at http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.html

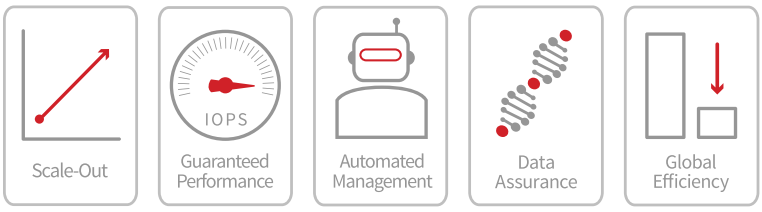

This paper introduces five architectural principles guiding the development of the next generation data center (NGDC). It describes key market influences leading a fundamental enterprise IT transformation and the technological trends that support it. The five principles are: scale-out, guaranteed performance, automated management, data assurance, and global efficiencies. Cloud infrastructure delivery models such as IaaS, private clouds, and software-defined data centers (SDDC) are foundations for the NGDC. In an era where IT is expected to ensure production-grade support with a plethoric flow of new applications and data, these models demonstrate how to eliminate bottlenecks, increase self-service, and move the business forward. The NGDC applies a software-defined everything (SDx) discipline in a traditional, hardware-centric business to gain business advantage.

These principles support a fundamental, integrated approach affecting the software, processes, and people that support each layer of the data center stack. They represent a paradigm shift from enabling technology at the speed of business to enabling business at the speed of technology. Driven by the unchained cloud economy offering innovations in infrastructure design and cloud services, the enterprise is well-positioned to embrace this shift. Cloud economics can be described as the discipline of maximizing network, compute, and storage resources in conjunction with cloud computing to satisfy customer needs. NGDC cloud economics offers a new discipline: one that uses IT resources as the glue for cost containment, the force for innovation, and the architect of new revenue streams.

> NGDC cloud economics: the discipline of using IT resources as the glue for cost containment, the force for innovation, and the architect of new revenue streams

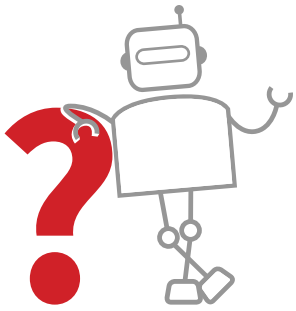

Chapter 1

Why Invest in the Next Generation Data Center?

Pressure from public cloud on IT delivery

Driven by higher expectations from internal and external market forces, the enterprise is taking notice. Public clouds like Amazon Web Service (AWS) have been demonstrating the possibilities for next generation IT deployment: rapid elasticity and high performance at lower costs.

Pressure from the cloud is evident in the shadow IT effect still simmering in the enterprise. The dissonance of providing rapid delivery while struggling with IT bottlenecks and the velocity of new applications created this substitute fulfillment process. A recent study in five key world markets found shadow IT, and its impact on IT budgets, continues to rise.[5] Shadow IT also lurks within the very department that seeks to eradicate it, as IT employees are more likely to adopt nonapproved SaaS than LOB users.[6] High-performing business units and high-performing IT developers have been going rogue, circumventing legacy infrastructure to satisfy their customers' needs.

In addition to its catalytic power to reshape enterprise IT service delivery, the cloud's promise of cost savings and on-demand service delivery has not gone unnoticed in U.S. government affairs, a traditional late adopter in technology. In 2012, the real-time video that was streamed so earthbound humans could watch the Curiosity rover land on Mars was made possible through Amazon's public cloud. NASA's Jet Propulsion Laboratory CTO Tom Soderstrom, explained, *"We realized eight years ago that cloud computing would be a game changer — we could do science in a few days that we would otherwise spend weeks on."* [7]

In the NGDC, developers spend less time on repetitive tasks and more time doing what they were hired to do: innovate and accelerate the business.

5. Atos Global Newsroom. "Out of the shadows: Stealth IT spend to increase by 20% in 2015." Available at http://atos.net/en-us/home/we-are/news/press-release/2015/pr-2015_03_26_01.html
6. Bloomburg.com. "McAfee Finds Eighty Percent of Employees Use Unapproved Apps at Work." Available at http://www.bloomberg.com/article/2013-12-04/a.dVwu6qQEIU.html
7. The Washington Post. "The rise of Amazon Web Services." Available at http://www.washingtonpost.com/business/on-it/the-rise-of-amazon-web-services/2015/04/23/fe1b3980-e795-11e4-9a6a-c1ab95a0600b_story.html

Automation enablement

The ability to cost-effectively deliver enterprise IT services on-demand and at scale is directly correlated to automation enablement. In traditional data centers, each infrastructure component has an independent set of management software and inherent integration between hardware and applications. Time-consuming processes, human error, and a lack of agility impede administrators managing legacy models. Dynamic workloads simply cannot scale in a traditional infrastructure without reconfiguration and hardware investments, putting more pressure on constrained IT time and budgets.

NGDC service offerings are abstracted from the underlying infrastructure. These services interface via APIs, translating the business requirements into resource offerings. Automation eliminates complex, rigid, hard coding to siloed projects. Resources are dialed up or down, and SLAs are rapidly modified at the application level to meet changing business-unit requirements. This control plane shift allows resources for compute, network, and storage to deploy faster and be managed easier.

Automated management enables the enterprise to capitalize on the agility and elasticity of the cloud. Staff changes will naturally evolve in an automated environment, as the ratio of servers to administrators increases. In the NGDC, developers spend less time on repetitive tasks and more time doing what they were hired to do: innovate and accelerate the business.

Innovation enablement

The enterprise knows that pooling resources offers exponentially more productivity and utilization; more performance and predictability; and less hardware, footprint, energy drain, and OPEX. IT isn't questioning whether to use new tools and methodologies. Instead, the questions become: How can we drive innovation that adds significant value? How can we seize opportunities to leverage the cloud economy and differentiate the business?

Innovation is at the heart of any CTO's mission. Today, enterprise IT groups have "a unique opportunity to move up the value chain," said Chuck Hollis, SVP of Converged Infrastructure Systems at Oracle. "The very nature and role of the corporate IT function is being questioned, evaluated, and debated in a way that we haven't seen in a very long time."[8] Every company, no matter the industry, is a de facto technology company today. Integrated into virtually every business process, IT competency is differentiating the enterprise more than ever in our mobile, on-demand marketplace. Enterprises adopting NGDC technologies consider data as business assets — yielding high value when data is innovatively applied.

Reduce complexity and risk

Enabling innovation without expanding risk is paramount for the enterprise. Legacy IT infrastructure is complex with multiple, disparate platforms and reliance on manual operations. Operational threats are heightened each time IT can't react quickly to customer needs. Building a cloud infrastructure requires highly organized design methodologies but eliminates the need for silos that elevate risk and hold IT back. NGDC architecture inherently minimizes risk with a converged approach to infrastructure management.

8. Chuck's Blog. "Hey Enterprise IT Vendors! Your Customers Need Help!" Available at http://chucksblog.typepad.com/chucks_blog/2013/10/hey-enterprise-it-vendors-your-customers-need-help.html

Chapter 2

The NGDC Evolution

Stage 1: standardize → virtualize

The development of new enterprise standards sought to reduce complexity and OPEX in the data center. The transformation applied to vendors, architecture, and IT approach. Today, the concept of standardization is still as compelling as it was over a century ago, when Henry Ford standardized industrial processes and reduced the time to build a car by sixfold.

In the last two decades, IT standardized on x86 servers and Ethernet-based networking. The challenges in managing infrastructure sprawl soon gave rise to the concept of consolidation. IT also standardized on employing server and desktop administrators. Managing sprawl likely involved consolidation of the underlying infrastructure and the people who support it, because both infrastructure and communication silos impact IT services and stifle the business.

Virtualization allowed enterprises to squeeze out more from what they had — gaining business agility, particularly in compute cycles. The data center was operating with larger, centralized resource pools. While silos still remained, networking was commoditized, and interoperability became the norm.

Stage 2: virtualize → automate

Virtualizing server resources allowed higher utilization, but innovators soon pursued faster, more efficient IT service delivery and turned to the cloud for solutions. The cloud extends all the virtues of virtualization by offering resources on-demand and as a service. In the cloud, IT administrators no longer control activities such as configuring, managing, and maintaining the hypervisor; these functions are administered through APIs that automate communication. Cloud computing allows the enterprise to maximize network access to a shared pool of elastic, scalable resources. TechTarget SearchDatacenter explains:

"Software-defined hardware was initially about consolidation — getting 20 physical servers virtualized onto three host servers running 20 virtual machines. The software-defined infrastructure of 2015 and beyond is more about rapid service delivery and flexible back-end system than utilization." [9]

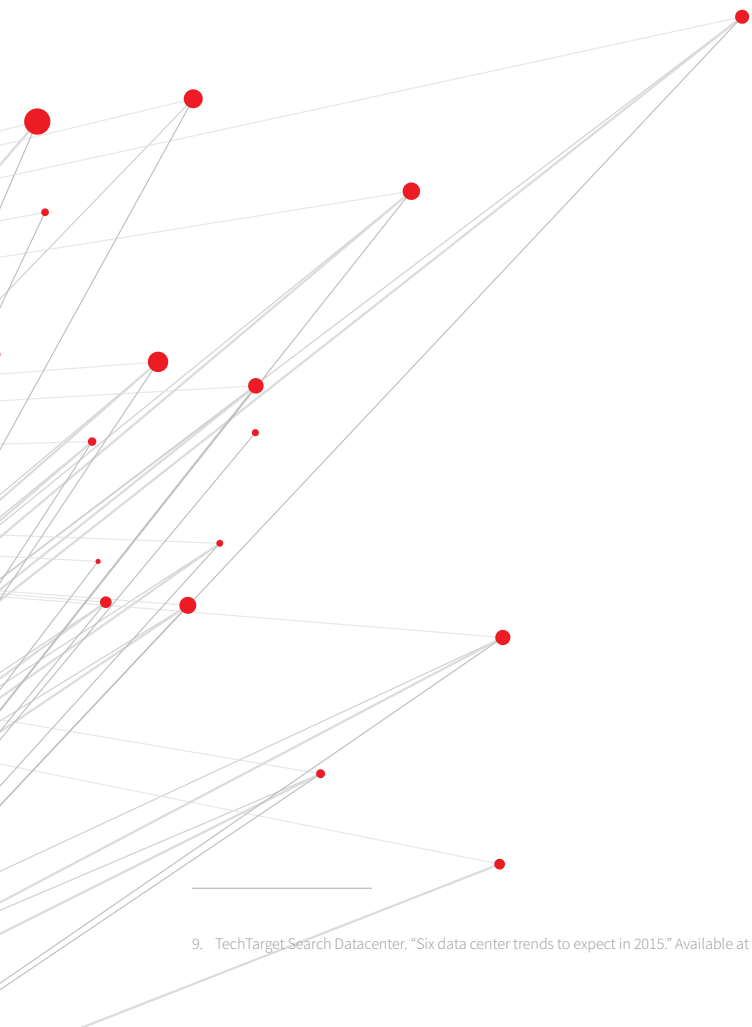

If automation is about codifying tasks, cloud orchestration is about codifying processes. By orchestrating workflows, delivery is automated in concert, including the arrangement, coordination, and management of complex systems, middleware, and services. Orchestrating means abstracting vendors into IT resources with services on-demand as the end goal. Orchestration software leverages all the advantages of automation by reusing the basic building blocks and stitching together the automated tasks into business-level workflows.

Think of the NGDC journey like eras of economic trading. A siloed, legacy IT infrastructure can be likened to everyone growing their own corn crop. In the virtualized era, their crop was pooled together; although corn was traded on the national market, it was still corn. In the automated era, the harvest is less specialized and more global; consumers don't necessarily know where their corn comes from, but it's always available and less expensive than before. In the NGDC, IT no longer has to concern itself with where the services are coming from; IT services are available on-demand, and expenses are reduced.

9. TechTarget Search Datacenter. "Six data center trends to expect in 2015." Available at http://searchdatacenter.techtarget.com/news/2240237521/Six-data-center-trends-to-expect-in-2015

Chapter 3

The Five Principles of a Successful NGDC

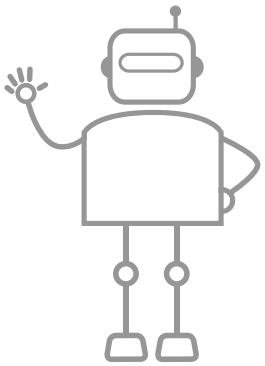

NGDC design involves a unified approach to applying five architectural principles across each layer of the data center framework. This approach affects the software, processes, and people supporting the entire delivery stack. The five principles are:

Scale-Out: In the NGDC, resource pooling provides nondisruptive horizontal expansion across the data center layers. Projecting business needs two years, three years, or five years out is a nonissue, because your infrastructure can scale incrementally as your requirements change. Linear, predictable growth of capacity and performance with guaranteed quality of service (QOS) is an essential operative in the NGDC, which makes a scale-out design critical to success.

Guaranteed Performance: Raw performance is only half the solution in the NGDC. QoS controls must be incorporated across the entire infrastructure or else any guarantee is only as good as the weakest link. Legacy infrastructures were not designed to balance increased network capacity simultaneously with expected levels of service in large-scale cloud deployment. In the NGDC, resources like CPU, memory, bandwidth, storage capacity, and storage performance are dynamically managed to deliver the application experience required and expected.

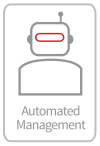

Automated Management: Enterprise IT is charged with enabling innovation and growth. If IT gets in the way, the enterprise risks revenue loss and significant disruption by more agile competitors. Employing software automation to take charge of decision points changes the paradigm from the physical limitations of hardware to the unbounded design capabilities of software. NGDC automation maximizes business results with policy-driven provisioning and resource allocation and eliminates endpoint-centric administration. Automating tasks and orchestrating workflows are fundamental in the enterprise if service-delivery needs are expected to be met at scale.

Data Assurance: NGDC data assurance includes seamless infrastructure resiliency without application configuration. Engineering for data assurance can be likened to designing buildings for earthquakes; there is no building that is guaranteed to withstand all earthquake activity, but earthquake-resistant structural design exponentially lessens the probability of disaster. NGDC architects plan for failure while mitigating its likelihood of occurrence with a self-healing, fault-tolerant architecture.

Global Efficiencies: The traditional data center model was built on the premise that more capabilities would require more resources — be it physical, financial, or human resources. In the NGDC, enterprises are not burdened with excess IT resources. Enabling global efficiencies begins with improved utilization of server platforms, networks, and storage protocols, as well as the vendors and IT teams that support each layer.

The evolution of enterprise IT infrastructure

Legacy data center		Next generation data center
Single tenant	>	Multitenant
Isolated workloads	>	Mixed workloads
Dedicated infrastructure	>	Shared infrastructure
Scale-up	>	Scale-out
Preprovisioned capacity	>	Capacity on-demand
Hardware-defined	>	Software-defined
Project-based	>	Self-service
Manual administration	>	Automation

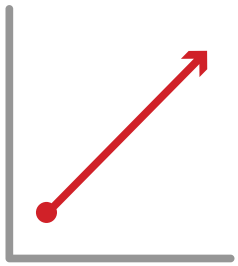

Chapter 4

Scale-Out

The foundational tenet of software control in the NGDC is represented in a scale-out design. Scale-out is positioned to become a predominant architecture in enterprise IT deployment because of its ability to offer seamless, transparent resource expansion without the cost and complexity of traditional infrastructure migrations.

Scale-up versus scale-out

Scale-up architectures allow expansion of a limited set of variables because the capability of each component limits their collective performance. Data and applications will usually proliferate and, ultimately, maximum capacity will be reached in a scale-up architecture. As applications are added, the infrastructure resources are spread out and performance degradation creeps in. With a scale-up architecture, the hardware you operate determines your limits, and faster hardware likely requires significant migration efforts.

By contrast, a scale-out architecture is a distributed architecture. It is not limited to the resource capacity of a single machine. It delivers linear expansion of all variables, allowing you to scale up or down by clustering resources. Scale-out offers a wide range of resource expansion options, allowing the enterprise to leverage IT investments consistently across the data center infrastructure and over the long-term.

The key to a successful scale-out architecture is making sure the performance from additional nodes can be intelligently provisioned by QoS capabilities. If workloads cannot be protected from each other in terms of performance, then application owners won't support consolidation. A scale-out architecture is designed for the enterprise where linear, incremental growth of IT resources is part of a global IT strategy.

Trends enabling scale-out

Resource clustering

The traditional, siloed data center relies on dedicated, physical servers with stranded resources. Those servers, along with their footprint, dictate the maximum data that can be managed or stored. Virtualization technologies such as VMware vSphere are built on the concept of clustering isolated resources into a shared pool.

This strategy is key for parallel processing applications, load balancing, and fault tolerance (redundancy) in the cloud. Pooling resources helps IT enable quick, nondisruptive modifications with singular control software. Adding another server to the pool effectively expands CPU and memory, and resources can be scaled in or out to meet changing business demands.

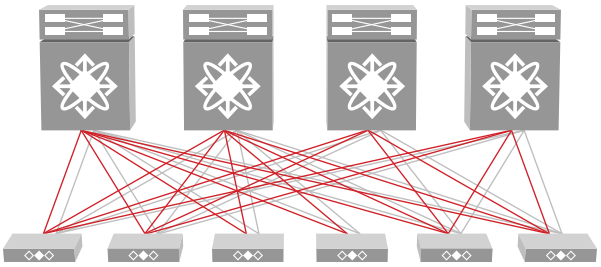

Leaf-spine

Before the adoption of Ethernet fabric, network topologies were built in layers that resembled a hierarchical structure and were typically segmented into pods that inhibited the location of VMs and other devices. Designed for use in enterprise networks, the traditional model consisted of core routers, aggregation switches, and access layer switches that traditionally required the use of Spanning Tree Protocols (STP) to facilitate scaling. This topology added significant complexity to both network design and troubleshooting.

Leaf-spine design has only two layers: leaf and spine. Access switches in the leaf layer connect to devices such as VMs, firewalls, edge routers, and load balancers. The backbone of the network is the spine layer with switches that perform routing. Each leaf switch is interconnected with every spine switch via dynamic Layer 3 routing and the use of Equal Cost Multi-Pathing (ECMP). Adjustments and determination of best path are based on responses to network change.

As leaf-spine architectures scale to thousands of nodes, latency and bottlenecks between access-layer switches – common in traditional architecture – are minimized because these switches are no more than a single hop away. Network vendors like Arista are enabling speed, density, predictability, and scalability with this NGDC topology.

Solid-state arrays

Even as compute and network administrators began leveraging the efficiencies of cloud, storage technology remained mostly static. In legacy storage, enterprises had little choice but to use separate storage networks and distinct pools for various performance levels and critical applications. Administrators managing traditional scale-up architectures were required to add a fixed number of drive shelves to add capacity. These disk arrays required upfront capital outlays for all resources anticipated for future expansion.

Scale-out, solid-state arrays are expanding the once siloed, static storage ecosystem. In the NGDC, pools of storage (GB, IOPS) are shared resources. The ability to cluster these resources provides incremental, as-needed, on-demand scaling; IT purchases and manages only what the business needs today. NGDC scale-out storage provides the flexibility to independently, nondisruptively scale both capacity and performance in a predictable linear pattern over time. This means the business can strategically scale in or out and distribute data and traffic over any number of nodes, while increasing the scope of data services. IT gains efficiencies by leveraging the cloud's elasticity in the storage layer and integrating tightly with all other services, resources, and cloud-management platforms.

Chapter 5

Guaranteed Performance

In the NGDC, raw performance is half the solution; just delivering service is not good enough. QoS resource controls must be utilized across the entire infrastructure or else any guarantee is only as good as the weakest link. CPU, memory, bandwidth, storage capacity, and storage performance is dynamically managed to deliver the application experience expected and required.

Trends enabling guaranteed performance

Monitoring tools

Information is power. Automated performance monitoring provides the critical intelligence needed to manage usage and availability of all IT resources. Enterprises are using software to monitor, report, and analyze live and historical data in the cloud. These tools are enabling service level improvements, managing security risks, and reducing operational costs.

Splunk Inc. is one of the companies forging a leadership position in monitoring tools. IT Service Intelligence (ITSI) is a monitoring and performance tracking platform for infrastructure analytics and visualization. ITSI runs in production environments with dashboards to map key performance indicators (KPIs) and infrastructure and provides drill-down data for problem evaluation. The software was developed for IT operations on premises, in the cloud, or in hybrid environments. Users cite incident investigation, stronger security, availability, and business-critical insights as benefits of implementing the software.

Network QoS

In the NGDC, enterprise architects are managing performance by segmenting network traffic into prioritized queues, while simultaneously controlling how much of the pipe is used for storage or applications. Network QoS is designed to guarantee service delivery for critical application traffic without breaking other flows, and to enable granular control of bandwidth utilization.

Applying QoS at the network layer is a key enabler to facilitating convergence and end-to-end service level agreements. Intel's Marketing Programs Manager Brian Yoshinaka explains the impact of network infrastructure choices:

"Today, as IT departments prepare to deploy internal cloud environments, it's significant to evaluate how network infrastructure choices will impact the cloud's ability to meet its service level agreements (SLAs). Terms commonly used to describe cloud-computing capabilities, such as agility, flexibility, and scalability, should absolutely apply to the underlying network as well." [10]

10. Data Center Knowledge. "Tips for simplifying your cloud network." Available at http://www.datacenterknowledge.com/archives/2011/12/14/tips-for-simplifying-your-cloud-network/

Storage QoS

In storage, the term QoS has lacked a consistent definition among industry vendors, creating understandable skepticism among IT decision-makers. Legacy storage models operate various virtual workloads on discrete pools of storage. The model was created to protect the performance and availability of a single application across a single storage array when each application has a diverse I/O pattern that must be supported. While these systems could provide capacity on-demand, they struggled with allocating performance resources efficiently, because they were not built to support the individual capacity and performance requirements of collective workloads. Enterprises then purchase more storage than they need, driving costs up and efficiencies down. ESG Lab Analyst Aviv Kaufmann explains the legacy platform challenge:

"Traditional storage infrastructures have evolved to better meet the demands of enterprise workloads by leveraging new technologies as they become available. But 'evolving' and 'incorporating' indicate a compromise; this is not the same as being purpose-built for the task. Traditional storage architectures can be configured to meet the SLAs of most of today's consolidated workloads, but to do so often requires overprovisioning and overspending for an infrastructure that is inflexible, complex to manage, and slow to respond to the needs of the business."[11]

In the NGDC, storage QoS means that performance controls are enforced at a granular level by ensuring a guaranteed amount of storage resources to each application. Every application provisioned is assigned a guaranteed amount of IOPS and those allocated IOPS are respected consistently, regardless of any other application activity, capacity level, or I/O pattern. The NGDC uses a node-based, shared-nothing storage design that can scale and consolidate workloads without increasing application risk profiles. It offers consistent, predictable performance based on the requirements of and value to the business, while isolating and protecting the workload from other workloads or hardware/software faults.

The precision balancing act between performance and capacity in NGDC storage design enables guaranteed performance with true QoS and reduced OPEX. What virtualization enabled in compute resources is now available in NGDC storage architecture: pooling and isolation of separate resources, reservation of resources for critical workloads, and moving resources dynamically from workload to workload.

> The precision balancing act between performance and capacity enables guaranteed performance and true QoS while reducing IT expenses.

11. ESG. "Quantifying the Economic Value of a SolidFire Deployment." Available at http://info.solidfire.com/rs/solidfire/images/ESG-Lab-WP-Quantifying-SolidFire-Value-Feb.pdf

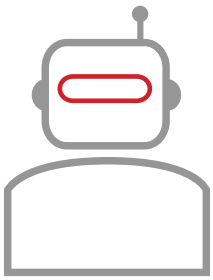

Chapter 6

Automated Management

The ability to consume resources on-demand and as a service is directly correlated to how deeply the enterprise can automate its infrastructure. Employing automation to take charge of decision points changes the paradigm from the physical limitations of hardware to the unbounded design capabilities of software. Automation across the stack is vital in the NGDC, where speed and innovation rule. Fast IT means faster innovation, faster time to market, and more value to the enterprise.

At Cisco Live in June 2015, John Chambers gave his final keynote as CEO of the networking market leader. In doing so, he chose to deliver a very direct warning to the audience of more than 25,000 attendees, some of whom may be slow to react to the disruptive digital renaissance. Chambers forecasted that 40% of businesses in the room would not exist in a "meaningful way" in the next 10 years.

In the face of his prediction, Chambers emphasized the role IT must play in successfully navigating this transition, stating: "IT no longer enables your business strategy; it IS your business strategy." Bluntly connecting the dots between IT strategy and business outcomes, Chambers suggested that "if you can't deliver on fast IT, your company will fail."[12]

Trends enabling automated management

The software-defined data center

Across new enterprise cloud consumption models, the software-defined data center (SDDC) can discover, deploy, manage, consume, release, and monitor infrastructure with fine-tuned control. Agile business systems require this approach to eliminate IT silos and shift from static to dynamic, fully configurable infrastructure operation. SDDC technologies accelerate the adoption of NGDC cloud consumption models by allowing a convergence of workload-centric architectures and cloud computing into a single domain.

In the NGDC, the control plane and data plane are loosely coupled or abstracted, and the control plane — which is making decisions about where traffic is sent — is implemented in software. IT has the ability to control, deploy, and reconfigure the infrastructure via software, specifically via open API control points. This has been referred to as the separation of control from data plane. When all infrastructure resources are abstracted into a higher-level control plane, they can be dynamically provisioned and managed, and resources are granularly defined to match application and/or service needs.

12. SolidFire blog. "Cisco Live! takeaway: Fast IT = Faster Innovation = Faster Revenue." Available at http://www.solidfire.com/blog/cisco-live-takeaway-fast-it-faster-innovation-faster-revenue/

Think of the control plane like a restaurant menu of options. You may order a complete meal or à la carte menu items. You choose what you want from the menu of available dishes (control plane) and tell your waiter. The data plane in this analogy is the kitchen, which cooks up your order as requested. You may not see how the chef is filling the order, but you appreciate a fresh and hot meal when it arrives at your table.

Taking this analogy a step further, consider new applications in restaurant automation that are improving service delivery and increasing revenue. At Chili's Grill & Bar, patrons can choose to self-order and self-pay via tabletop tablets. Chili's cites a jump in sales as customers no longer have to flag down wait staff to refill a drink or to order dessert. In addition to increasing revenues, the company sees this as an opportunity to create stronger connections and interactivity between the restaurant and its guests.[13]

Application programming interface (API)

Today's API ecosystem is diverse and expansive. IT operations leverage multiple tools to achieve diverse goals. Offering CLIs can be insufficient; vendors are illustrating the value and benefit of code as an instrument for day-to-day operations. Scripting common and/or disparate infrastructure tasks with APIs dramatically reduces errors, increases productivity, and streamlines service delivery by enabling efficient sharing of data and processes. Speed and agility are fueled by repeatable methods that validate and verify tasks. A consistent and repeatable process provides an auditable record of activity and a gateway to fully automated management.

Employing software automation to take charge of decision points changes the paradigm from the physical limitations of hardware to the unbounded design capabilities of software.

13. Computerworld. "Automation arrives at restaurants (but don't blame rising minimum wages)." Available at http://www.computerworld.com/article/2837810/automation-arrives-at-restaurants-but-dont-blame-rising-minimum-wages.html

Configuration management tools

In the last several years, the automation ecosystem has seen explosive use of configuration management (CM) tools to automate DevOps processes and actions across the infrastructure. These tools help orchestrate elastic workloads in the cloud at scale. Configuring and maintaining from a few servers to thousands is easier, because the infrastructure can be described and managed with code. A flexible, policy-driven approach to IT operations enables enterprise cloud agility. In the NGDC, IT administrators manage the full spectrum of enterprise solutions via policies that govern behavior.

Puppet was conceived as a cross-platform, open-source CM solution. In 2011, with investments from partners such as VMware, Google, and Cisco, Puppet Labs introduced its first solution for commercial enterprise. Its "declarative configuration language" helps define and manage each step in the delivery process and infrastructure lifecycle, including provisioning, configuration, OS management, orchestration, and reporting.

Chef offers a similar CM tool, also based on Ruby, for both open source and enterprise CM solutions. With Chef, a workstation is deployed to control the master server. Its configuration uses Git, an open-source distributed version control system. Chef Director of Patterns and Practices Colin Campbell explains the product strategy: "Treating infrastructure as code produces a workflow that effectively can use the same testing process used for other types of software."[14]

Software-defined networking

In conventional networks, operators configure functionality at a low level, device by device. With software-defined networking (SDN), administrators can write high-level control programs and stipulate the behavior of an entire network, because the control logic is separated from underlying physical routers and switches that forward traffic. Complex tasks for integrating network functions such as resource control, prioritization, and security are specified into a single SDN control framework. Paradoxically, SDN offers the ability to create configurations that have more sophisticated policies but are much easier to manage, maintain, and secure.

Announced at VMworld in 2013, VMware's NSX was an early entrant in SDN implementation. Born out of VMware's acquisition of Nicira in 2012, NSX is a network virtualization platform that reproduces the network model according to policies set in software, enabling any size network topology to be rapidly created and provisioned. Built around a controller cluster that manages the distribution of logical network functions (switches, routers, firewalls, VPN, security) into hypervisors, NSX allows virtual networking through any platform leveraging NSX APIs.

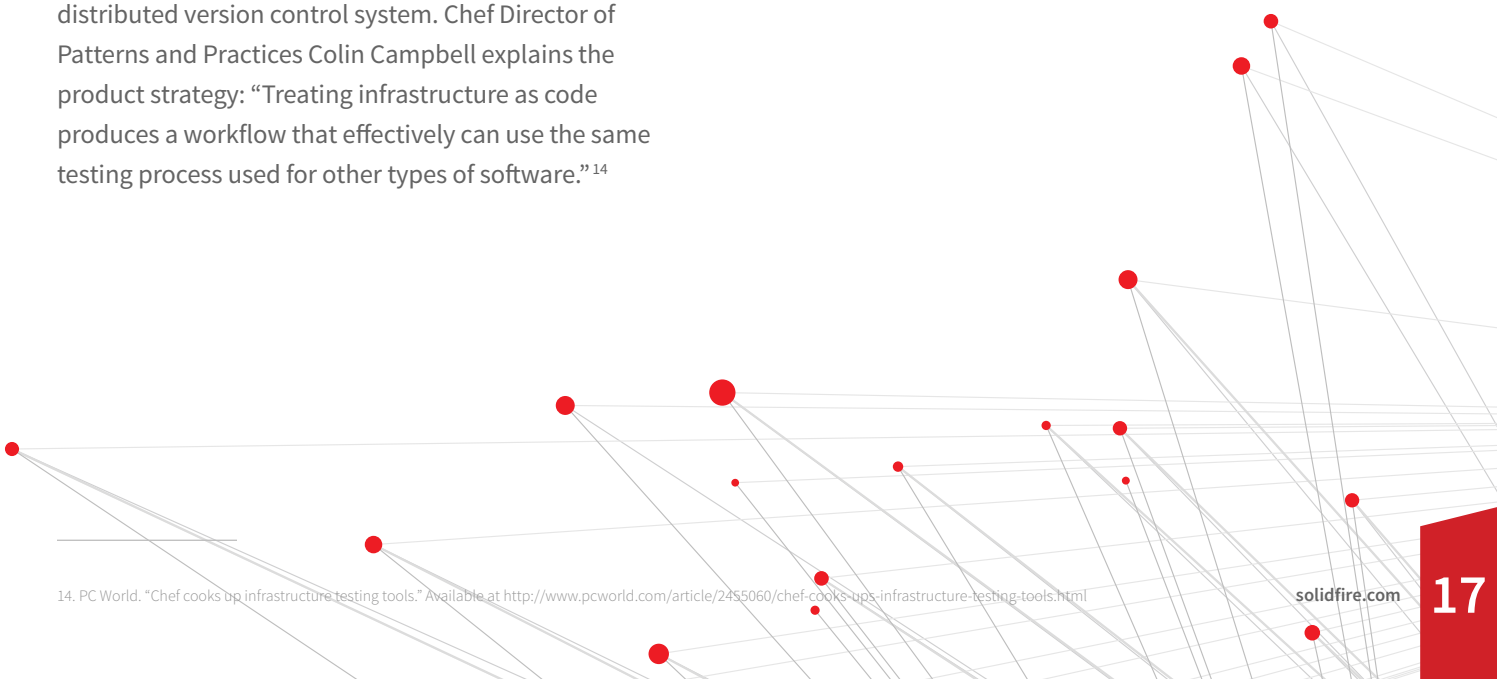

14. PC World. "Chef cooks up infrastructure testing tools." Available at http://www.pcworld.com/article/2455060/chef-cooks-ups-infrastructure-testing-tools.html

This approach helps IT use the physical network as a pool of capacity for on-demand consumption. Chris King, Vice President of Product Marketing for Networking and Security at VMware, says businesses are demanding a change: "The transformation of the network is inevitable. The only real question is: Why would you entertain an antiquated architecture when so much of your business's success relies on your ability to deliver more speed and efficiency? Why should agility require a forklift? Virtualizing your network is the next step forward on the path to the software-defined data center."[15]

Another example of this trend is Cisco's ACI framework. ACI is an integrated network architecture for dynamic workloads that is moving the needle in network automation. The ACI automated application-centric policy model promises to meet IT market needs with embedded security, centralized management, compliance, and the ability to quickly scale.

ACI is reported to be the industry's first solution to provide a dynamic, application-aware network policy that can reduce application deployment time from months to minutes.[16] Cisco reported in June 2015 that ACI user and leading security-solutions firm Symantec projected a five-year cumulative benefit of 441% ROI and reported 87% faster application development lifecycle as a result of deploying ACI.[17]

Storage APIs and interfaces

In storage, automated management is likely to be a combination of interfaces to manage the system directly (CLIs and UIs), integrations to orchestrate higher-level tools (SDKs, plug-ins, drivers), and storage APIs to automate workflows (adding, deleting, or reporting on usage). The current paradigm treats storage as an undifferentiated resource in the IT automation puzzle. Applications and workloads in the cloud need more than just capacity-provisioned storage; they need profiles driven by business requirements, rather than the limited options offered with legacy storage architectures.

Storage automation in the NGDC changes the capacity-provisioned paradigm to an automated, programmable, repeatable, policy-driven environment that drives enterprise IT transformation. It eliminates device-and endpoint-centric administration and fully enables multitenant data management and protection. Web-scale principles of automation, QoS, and API-based access help orchestrate all enterprise workloads and their integration with NGDC delivery solutions. These attributes enable IT to deliver much higher business value for the enterprise.

15. Network World. "Virtualize It! Build a data center that is defined by software not by hardware." Available at: http://www.networkworld.com/article/2359521/tech-debates/what-s-the-best-approach-to-building-next-generation-data-center-networks.html

16. Network World. "Cisco's ACI goes beyond SDN with policy-based automation." Available at http://www.networkworld.com/article/2359521/tech-debates/what-s-the-best-approach-to-building-next-generation-data-center-networks.html

17. Cisco. "Cisco Extends SDN Leadership; Delivers Customer Choice for Data Center Programmability and Automation; Enhnaces Software and Expands Nexus Portfolio." Available at http://newsroom.cisco.com/press-release-content?articleId=1648780

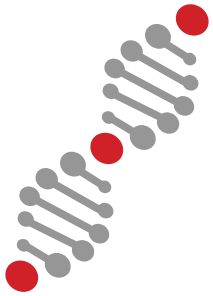

Chapter 7

Data Assurance

In the legacy data center, infrastructure was customized with specialty hardware to ensure a failure would not negatively impact specific applications. The applications were saved, and the price was paid in over-engineering.

Engineering NGDC data assurance can be likened to designing buildings for earthquakes; there is no building that is guaranteed to withstand all earthquake activity, but earthquake-resistant structural design exponentially lessens the probability of disaster. In the NGDC, a resilient and secure infrastructure is the expectation, and failure is managed because the self-healing architecture is designed to tolerate it. From component failure to failures in managing business unit dynamics, NGDC enterprise architects plan for failure while mitigating its likelihood of occurrence.

In the NGDC, a resilient and secure infrastructure is the expectation, and failure is managed because the self-healing architecture is designed to tolerate it.

Trends enabling data assurance

Dynamic CPU and VM allocation

VMware's DRS (Distributed Resource Scheduler) is a utility that dynamically allocates and balances computing capacity and VM placement with pooled resources from multiple ESX server hosts. Using automatically programmed, predefined priorities, hypervisors can redistribute VMs from areas of contention to other areas without workload disruption. Enterprises are using DRS to simplify provisioning of applications and to get higher utilization by optimizing resource allocation.

The business benefits of utilities and features such as DRS include the ability to simplify day-to-day IT operations, as staff members are less affected by localized events in their environment. When loads on individual VMs change, automatic resource optimization and relocation will reduce the need for administrators to respond. Time is dedicated to higher-level tasks and risk of data loss or service disruption is avoided.

Availability zones

The concept of availability zones for data assurance in the cloud was popularized by AWS. Availability zones protect your applications from the high cost of failure. In open-source platforms such as OpenStack, availability zones are used to arrange compute hosts into logical groups, providing a form of physical redundancy and isolation from other availability zones. Users can provision resources by specifying which availability zone they want their instance to be built in, allowing application resources to be spread across disparate machines for high availability (HA) in the event of a failure.

Multi-pathing

Similar to a utility power grid that provides multiple paths to a single location, transparent multi-pathing allows data to pass via alternate routes. This removes the likelihood — and risk — of a single point of network failure, whether its component or environmental. In addition to enabling redundancy, NGDC technologies like multi-path routing are improving network performance because more load-balanced traffic can be distributed through more routes.

Enterprise architects are using multi-pathing to gain higher throughput, reliability, and performance in the network.

Dynamic firewalling

Another networking trend in NGDC deployment is dynamic firewalling to segment-specific instances or groups of instances. Administrators programmatically apply firewall policy to a workload or group of workloads instead of manually touching each device in the chain, which can be upwards of 50 endpoints in a large-scale cloud deployment. Automating with dynamic firewalling enables data assurance by reducing the likelihood of human error and security breaches.

Replication-based HA in storage

Storage innovations have evolved in the last several years to ensure high availability, security, and data protection in the cloud. Enterprises that need to scale business-critical applications with guaranteed levels of performance are replacing RAID data protection with replication-based HA. A 30-year-old technology, RAID has not kept pace with the scale of cloud. Recovering from a failure with a traditional RAID-based storage design can take hours or days, risking greater faults and complete data loss.

Legacy storage companies built their management systems in the construct of using storage administrators, and, as a result, have created complex, feature-rich administration tools. Data assurance in the cloud is not realized through armies of administrators, but through effective configuration management tools built for the cloud. In the NGDC, enterprises are not administering their storage, they are automating it.

Replication-based HA is a post-RAID data protection scheme based on a distributed-replication algorithm. This NGDC architecture protects against both drive and node failure. It allows the storage system to absorb concurrent failures across all levels of the storage solution. In a failure event, each drive in the system redistributes a small percentage of its data in parallel to the free space on all other remaining drives, and no operator intervention is required. The self-healing properties of an automated, replication-based architecture enables fast rebuilds with minimal performance impact, because parity reads are not required. With this architecture, faults are isolated, and QoS settings remain enforced.

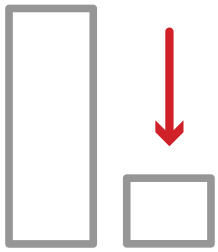

Chapter 8

Global Efficiencies

IT decision-makers are consolidating and converging the people, software, and processes that fueled under-utilized silos and IT sprawl. Consolidation enables IT to take charge of perceived activity, seasonality, and typical usage, while controlling oversubscribed resources. Like airline carriers that practice overbooking, administrators can programmatically govern thresholds to meet customer expectations, maximize utilization, and control costs. Fewer systems and fewer interfaces between components are pillars of global efficiencies in the NGDC.

Trends enabling global efficiencies

Converged networking

The cost/benefit of converged networking is gaining more attention from enterprise IT architects seeking optimization of dynamic NGDC services at scale. Rapid increases in Ethernet speeds, coupled with network-based QoS technologies, facilitate the convergence of previously separated application and storage traffic. Network designs can now leverage the cloud's elasticity. Further efficiencies are gained by eliminating Fibre Channel host bus adapters and per-port costs.

Network convergence helps enterprises leverage the agility of cloud inside their firewall, with the same web-scale technologies that power the largest public cloud providers. In the NGDC, a single-fabric technology reduces infrastructure and power requirements, optimizes network performance and utilization, and simplifies management and operations.

Native in-line storage efficiencies

A converged approach is also applied in NGDC storage architecture. Consolidation of storage into a simplified, scalable platform improves utilization rates and can substantially lower capital cost. The impact of all-flash storage on data center efficiencies cannot be overstated. All-flash reduces storage equipment costs, carbon footprint, and power consumption. NGDC data reduction techniques enable devices to store data using fewer bits, requiring less space. Architectural features such as in-line deduplication, compression, and data thin provisioning enable environmental and economic benefits at scale without performance trade-offs.

Conclusion

Enterprise architects are charged with improving TCO, meeting SLAs, and creating an environment that supports better, faster, and easier execution, all while ensuring production-grade support with a plethoric flow of new data and applications. Hardware-defined, legacy infrastructure constrained IT from achieving these goals — and understandably so. The traditional data center was not built for the cloud.

The new enterprise IT isn't chasing the curve; it is the curve. Service delivery models such as IaaS, private cloud, and SDDC will continue to disrupt traditional infrastructure approaches. NGDC cloud economics is changing the paradigm. Applying a software-defined everything (SDx) discipline in a once hardware-centric business creates more opportunity to maximize resources and satisfy 21st century customer needs.

The traditional data center was not built for the cloud.

SOLIDFIRE
Now Part of NetApp

solidfire.com
1-855-376-4752
info@solidfire.com