INTEL SDI ENABLES INTERNET OF THINGS (IOT) INTELLIGENCE

SCALABLE DATA MANAGEMENT AND BIG DATA EXPAND USAGE MODELS FOR IOT

EXECUTIVE SUMMARY

Intel's concept of software-defined infrastructure (SDI) extends the definition of the software-defined datacenter (SDDC). SDI is a re-evaluation of system architecture driven by the requirements of business flow, workloads, and specific applications—not by a menu of hardware available to purchase at the moment. SDI could transform mainstream datacenters and has the potential to displace current datacenter infrastructure and highly available transaction processing systems by the end of this decade. Until now, SDI has been conceptual, but Intel is working to enable real-world usage models to turn the SDI vision into a reality.

At the same time there has been a rapidly growing interest in the Internet of Things (IoT) driven by significant reductions in the cost of sensors, bandwidth, and computing. In December 2014, Intel announced the Intel IoT Platform: an end-to-end reference model that utilizes products from both Intel and industry partners. This paper provides an overview of SDI as it relates to data analytics in an IoT environment and highlights vertical use cases in manufacturing and retail as examples of how SDI can help optimize an end-to-end IoT value chain.

The next step to allow cloud datacenters to take advantage of SDI will be the creation of a defined set of use cases and reference architectures where SDI and other "software-defined" technologies can provide significant benefits. MI&S believes that data analytics driven by the Internet of Things (IoT) value chain could be an effective use case.

INTEL SDI: THE FUTURE OF CLOUD ORCHESTRATION & MANAGEMENT

MI&S believes that IT is on the cusp of a major datacenter architecture transition. This transition will be driven by 24×7 global business reach, dramatically increased use and depth of business intelligence (BI) and predictive analytics (Big Data), and pushing sensors and intelligence into our physical world in the form of the "Internet of Things" (from datacenters to wearable consumer electronics). It is impossible to predict exact technology directions even in a three-to-five year timeframe, but the industry is starting to form a good, high-level framework for the future of IT operations.

Software-defined datacenters (SDDC) are the next generation operational target for IT. SDDC is the flexible, agile infrastructure framework that will promote IT to a full business partner capable of creating value and increasing operational efficiencies. Though the term SDDC was coined by VMware, we believe the key concepts of SDDC are broader. Specialized hardware will have a role in future datacenter architecture, but infrastructure software will mediate workload access flexibly to scalable hardware.

Over the last 18 months, Intel introduced the concept of software-defined infrastructure (SDI) which extends this definition of SDDC. Intel's vision for SDI is to allow a user to define a set of characteristics that an application requires (key value store, persistent storage, *etc*.) and allow the infrastructure to respond by assigning an optimal set of available resources—without the user requiring specific knowledge about the underlying hardware. These resources are adjusted dynamically over time to meet scale and service level agreements (SLAs). Intel SDI defines three fundamental operational layers (Figure 1).

FIGURE 1: OPERATIONAL LAYERS FOR INTEL SDI FRAMEWORK

Orchestration

- Mid-level system policy engine software
- Enables higher level frameworks to dynamically manage service composition without interrupting ongoing service operations

Composition

- Low-level layer of system configuration software
- Continually and automatically manage a pool of hardware resources without interrupting ongoing operations

Hardware Pool

- Composable pool of modular hardware resources
- Independent compute, storage and networking hardware capabilities with standardized abstracted software interfaces

Intel is working with key partners and the open source community to drive their SDI vision into various IT focus areas such as OpenStack, general hosting, cloud services, Big Data analytics, telecommunications, and enterprise workloads. However, the next step to allow cloud datacenters to take advantage of SDI will be the creation of a **defined set of use cases and reference architectures** where SDI and other "software-defined" technologies can provide significant benefits. MI&S believes the data analytics driven by the Internet of Things (IoT) value chain will be an effective use case. Cloud datacenters should consider data analytics when looking to implement SDDC practices that encompass SDI, SDN (software-defined networking), and NFV (network functions virtualization).

THE INTEL IOT PLATFORM & DATA ANALYTICS

In December 2014, Intel announced the Intel IoT Platform: an end-to-end reference architecture that utilizes products from both Intel and industry partners. The goal is threefold…

1. Provide a repeatable foundation for securely connecting devices,
2. Deliver trusted data to the cloud, and
3. Provide value through end-to-end analytics.

Figure 2 provides an end-to-end view of Intel's IoT Platform with a focus on the IoT Analytics high-level architecture.

FIGURE 2: INTEL IOT PLATFORM

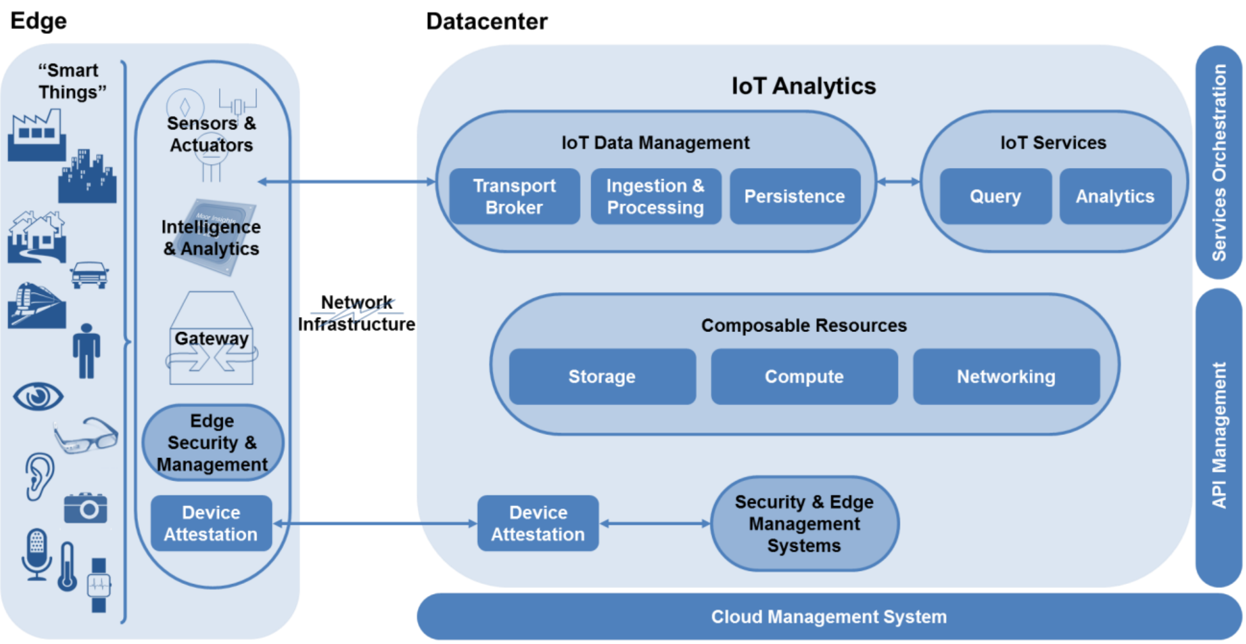

SDI'S ROLE IN THE IOT VALUE CHAIN

Figure 3 illustrates how SDI can be applied to Intel's current IoT platform. In this case, SDI can rapidly allocate, provision, configure, and deploy services on pools of hardware and software infrastructure resources across the distributed endpoints, network infrastructure, and cloud datacenter.

FIGURE 3: SDI IN THE IOT VALUE CHAIN

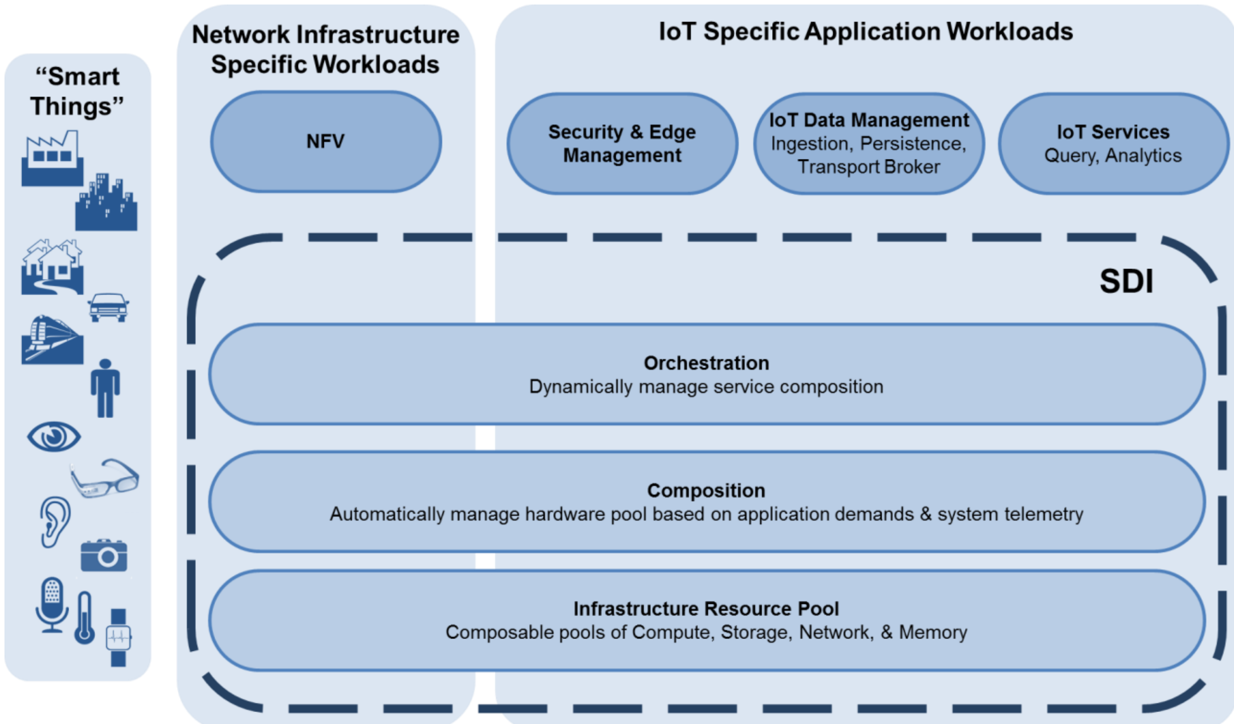

The rise of the Internet of Things is driving exponential increases in the amount of data that flows from IoT endpoints all the way through to the cloud—a "data deluge". It all begins with the sensor or actuator at the endpoint, and SDI offers the increased levels of efficiency to address the dynamic nature of IoT Analytics (both periodic and event based data generation). The types of data collected and technology used may be different for each use case. Regardless of the application, there is an increasing desire for more intelligent endpoints, aggregation capability across multiple endpoints, and analytics/processing horsepower at the edge of the wireless network to pre-filter data and provide real-time insights and actions.

For example, to increase intelligence at the "edge", network service providers are investing in initiatives to enable intelligent cloud-hosted nodes to be deployed into base station cell sites that also serve as IoT gateways. These nodes serve as distributed analytics engines to perform local processing for low-latency, location-aware insights in real-time—which provides the ability for network service providers to monetize services for this data while it is most valuable. Organizations such as the Mobile-Edge Computing (MEC) Industry Specification Group (ISG) have been formed to create standards for this purpose.

AGILE, FLEXIBLE NETWORKS

IoT data drives a surging demand for connectivity throughout the entire value chain. Service provider networks that historically have been built to serve predictable and measurable traffic profiles now are entering a world of IoT where network demands are rapidly evolving and not well understood. The ability for service providers to plan and deploy networks carefully in support of known demand is challenged by the variety of endpoints, the diverse set of applications, the nature of data, and the direction of traffic.

With this challenge in mind, service providers are looking to build agile, flexible networks with the ability to deploy and repurpose resources as demand changes. In addition, service providers are investing heavily in open-standards-based equipment without vendor dependencies to ensure they have maximum flexibility for ingesting data sources from as many points as possible. Software-defined networking (SDN) and network functions virtualization (NFV) strategies will allow service providers more easily connect the endpoint (sensor) to the backend analytics platform. In this case, SDN and NFV help the service providers morph and adapt quickly to changing needs, repurpose infrastructure with agility, and deploy their control and data plane flexible based on the cloud and networking technologies.

SMARTER NETWORKS

In addition to improving network flexibility and agility, service providers are looking to make their networks “smarter” by analyzing all of the customer data flowing through the network. With this data, service providers look to understand and predict customer behavior in real-time. To achieve this goal, service providers are investing heavily to make sure they can access data sources at every point in their networks and feed this data into their own analytics platforms. Leveraging this data and analytics capability could provide a new revenue source for service providers: by exposing the network data via application programming interfaces (APIs) to third parties who can leverage the data to provide improved intelligence into their services. This model is somewhat analogous to the Google model that gives a high value service away for free, and then monetizes the intelligence derived from the data. This capability opens up a whole new world of possibilities for service providers to enable new business models.

Once the data makes its way through various aggregation points, gateways, and network infrastructure, the cloud datacenter is responsible for providing seamless data ingestion capabilities by managing data from various streams. As the number of data sources continues to grow, data normalization across a broad spectrum of sensors, protocols, and closed loop control systems to ensure security and data accuracy all become increasingly important. The cloud datacenter is also responsible for ensuring that edge devices and gateways can be set up and securely discovered in seconds. And with the wide range of device types and services throughout the IoT value chain, it is critical to have a robust set of APIs to account for various protocols, service delivery software applications, and device types. Once the data is aggregated, the cloud datacenter uses Big Data analytics to provide more complex intelligence and insights.

IoT Frontend

A cloud datacenter that supports an IoT infrastructure requires increased levels of scalability, agility, and multi-tenancy capability. To support these requirements, Intel's IoT platform has two major datacenter front-end components.

- **Security and Edge Management Systems**
  - Ensures that endpoints sending data are who they are supposed to be and determines the level of service and security priority their data will have
  - Ensures that endpoints receiving data know the access and permission levels of the services communicating with them
  - Ensures that endpoints are properly managed: provision, update, configure, *etc*.
- **Data Management**
  - Resources are assigned to incoming (ingest) and outgoing (egress) streams as per security and service permissions.
  - The adjacent table describes what each stage of data management does.

Data Management

	Transport Broker	Ingestion & Processing	Persistence
Ingest (Incoming stream)	Identifies incoming streams from endpoints & services, then routes to optimal processing resources	Decrypts & decompresses incoming data, plus service-specific stream filtering, data normalization, and load balancing	Applies service rules & routes streams to appropriate storage/query or compute/analytics resources
Egress (Outgoing stream)	Creates & inserts downstream routing information for optimal service delivery	Orders, prioritizes, compresses, & encrypts data for most efficient transmission & endpoint ingest	Pulls data streams out of storage/query & compute/ analytics tasks as requested by remote services & endpoints

Concurrency is inherent in this model. Incoming data (ingest) can be sent directly to both storage and to compute resources—and therefore to both query and to analytics services—one, the other, or both at the same time. Outgoing data (egress) can be merged from many storage and compute resources.

COMPOSABLE RESOURCES

Underneath the frontend components are the typical compute, storage, and network resources found in most modern datacenters. However, under Intel's SDI architecture these components are designed to be **composable**. They could be configured dynamically to build and scale applications that deliver a managed quality of service (QoS) to higher level software services.

IoT data management could be built on the SDI pool, as well as the back-end, low level query and analytics services in Intel's IoT Analytics platform. It is important to note that IoT is in an early phase of maturity. Software applications that are functionally similar may have strikingly dissimilar requirements for compute, storage, and network provisioning. And those applications are evolving rapidly, so their requirements are likely to change anyway. **Composable pools of resources will enable the most efficient use of resources over time and across these emerging application components.**

Higher-level services will take advantage of query and analytics via services orchestration. Intel will expose APIs and API management tools to those building IoT services, so they can build their own capabilities. Intel continues to make configuration and detailed telemetry from server infrastructure based on Intel processors and other products available to partners cloud management systems. For example, Intel is building a reference data-as-a-service (DaaS) capability on top of their query and analytics APIs.

Composable Resources

Composability is enabled at a hardware subsystem and even at a component level. Each component in a system exposes information about its features, functions, performance, *etc*. (called "metadata") via an Application Programming Interface (API). The host OS or hypervisor tracks all of the hardware resources as pools of allocated and unallocated hardware resources. When a new instance of an application or service requests resources to run, resources are allocated based on availability metadata.

Subsystems and components do not have to be in the same physical location to participate in a composable pool of system resources. SDI looks at "systems" as logical constructs spanning racks, rows, and datacenters. Intel will push composability out to the edge of networks.

Intel's long-term vision for SDI and IoT is software-defined. SDI will use the control plane to gather network and IoT endpoint metadata. The control plane will manage policies for how data is analyzed at endpoints, in the network, or in a datacenter. Policy decisions for moving and analyzing data will be based on real-time metadata and statistical analysis, based on networking latencies and processing capabilities available in potential data paths between endpoints and datacenters.

Analytics Throughout the Value Chain

Intel's SDI vision could enable many usage models to drive analytics throughout the entire IoT value chain.

- **Data Management** Intel SDI could allow cloud datacenters to compose the right pools of resources for different kinds of analytics, backend services with multiple devices, and management and control of different sensor inputs.
- **Services Frameworks** Intel has partnered with Cloudera and the open source community to understand requirements and create accelerated solutions to deploy Hadoop and Spark analytics on top of both SQL and NoSQL databases.
- **Edge Analytics** enables intelligent gateways—and in some applications even intelligent distributed endpoints—that include processing and aggregation capabilities closer to the data.
- **Cloud Management** Intel SDI provides for a horizontal cloud management framework and policies to fulfill specific workload/tenant requirements for multi-tenant cloud datacenters.
- **Security and Edge Management** Intel SDI framework can help manage intelligent "micro clouds" at the edge by providing acceleration, aggregation, control, and policy management capabilities.

Use Case 1: Enabling the Smart Factory of the Future

The manufacturing sector is a very data-driven industry. Even the simplest factories depend on data to track production runs, quality metrics, yield results, tool/machine health, and inventory. However, often much of the data collected on the manufacturing floor is saved for a rainy day and not used proactively to generate business insights or improvements. In a world of razor thin margins, every penny counts: deploying new systems and tools to run their lines more efficiently could mean big savings to the bottom line. However, these razor thin margins also mean investments in new tools, processes, and techniques all must be evaluated carefully to ensure maximum ROI.

Leading edge manufacturing organizations are looking at ways to make their factories smarter. This involves finding more efficient ways to turn descriptive analytics (what happened and why) into predictive and prescriptive analytics (what will happen and what is the best course of action). Aggregating data from numerous sources provides more opportunity to achieve industrial intelligence. **Analyzing data at the source provides the lowest latency and highest responsiveness**.

Big Data analytics and the Internet of Things in manufacturing as an end-to-end platform is the key element that will enable smart manufacturing. In many cases, it is not about creating additional sensors/endpoints within the factory floor, but rather collecting and aggregating data in a meaningful way to translate it into insights.

In a smart factory environment, distributed gateways local to the factory floor can allow for real-time analytics for time-sensitive decisions. IoT gateways allow factories to connect tools, devices, and sensors to a central datacenter; merge and correlate a wide range of data sources and types; and use Big Data techniques to run their manufacturing more efficiently. Furthermore, ongoing use of data analytics can improve product quality, increase manufacturing throughput, identify root causes of manufacturing issues, and reduce machine failure and downtime.

Many segments within the manufacturing sector are not leading edge adopters of new technology. To spread wider use of “smart factory” practices, it is important for industry leaders to create a set of solutions guides and reference implementations for how to pilot and deploy end-to-end IoT to datacenter analytics capabilities.

The SDI vision as it applies to the smart factory will encompass everything from real-time analytics capability on the factory floor all the way through to Big Data analytics in the datacenter. On the factory floor, SDI can provide management and orchestration layers to ensure distributed gateways and endpoints can be discovered, deployed, and managed securely and consistently and also communicate with the cloud regardless of the application, protocols used, or device location. Testers, manufacturing tools, and automation databases are generating data in larger volumes, with more variety, and at a faster rate than ever before. An edge computing gateway local to the factory floor that uses SDI to orchestrate and manage the data across the endpoints could be used to efficiently handle real-time analytics and to filter the vast amounts of data. Thus only the most high value and relevant data is consolidated and sent to the datacenter.

SDI Benefits

As it moves through the network through various gateways, SDI can help normalize the data, maintain communications protocols, provide composable pools of network resources to optimize data flow between factory and datacenter, and improve interoperability.

In the datacenter, SDI practices can provide consistent scalable Big Data analytics that can be used across all manufacturing data types and use cases.

Intel’s industry leadership in manufacturing makes them well-positioned to help drive reference implementations for broader adoption. Intel conducted an IoT pilot deployment on one of their own factory floors. It shows how data analytics, applied to factory equipment and sensors, can be used to improve equipment component uptime, minimize wrong classification of good units as bad, enable predictive equipment maintenance, and reduce component failures. The pilot involved a solution developed in partnership with Cloudera, Dell, Mitsubishi Electric, and Revolution Analytics (which was acquired by Microsoft in January 2015).

USE CASE 2: INVENTORY OPTIMIZATION IN RETAIL

One of the key foundations for a successful retailer is to provide a great experience for shoppers wherever they are and whichever channel they use (m-commerce, e-commerce, physical store, *etc*.). Inventory optimization is fundamental to a positive customer experience: providing what customers want in the time and place they want it while maintaining a defined set of operational goals and cost structures. A “mom and pop shop” knows what to stock based on personal interactions with their customers and what their inventory looks like based on the small physical size of their store. Online retailers use clickstream data to personalize shopping experiences and can promote items where physical inventory is available. Larger retailers traditionally have lagged behind with limited visibility into what is stocked on the store floor and inventory distortion challenges (out of stock, overstock) which disrupt sales and profitability.

Various technologies including RFID, video surveillance, point-of-sale, and in store kiosks are available to provide additional visibility into what is happening on the retail floor. To track each item individually, RFID can make business sense for higher end goods (apparel, electronics, *etc*.) due to the cost of the RFID infrastructure.

Today, much of the data collected from retail endpoint devices is neither used nor understood, so it provides limited value for a retailer to improve the customer’s experience. Leading retailers are looking at ways to increase value through IoT technologies combined with data analytics tools. This combination can be used to trigger real-time actions (*e.g.*, restocking shelves when inventory is low, opening a new checkout lane when lines get long) **and** create longer-term insights (*e.g.*, product popularity based on “clickstream” data through store, optimizing store flow).

Video monitoring may be a better solution for retailers selling a mix of less expensive items like groceries. Video monitoring is not being used broadly today, because the capability is only now becoming possible via advances in machine learning, pattern recognition, low cost and high resolution sensors, *etc*. Intel is working with retailers to prototype and commercialize vision-based store shelf monitoring systems to further promote adoption.

The Intel IoT platform, when applied to retail as a reference architecture, makes it more cost-effective and easier for retailers to implement inventory optimization practices and combine them with Big Data analytics tools. At retail endpoints, large amounts of data on browsing history, inventory placement, and buyer behavior are being captured that can help retailers improve the customer experience via improved inventory placement, store layout enhancements, and adjacency selling opportunities.

Retailers are looking to create more responsive stores where the store itself composes its computing resources. Compute happens locally to provide real-time analytics capability, device orchestration, and compression and filtering of data before it is sent to the enterprise for more comprehensive data analytics. SDI practices will be a critical element to optimize these distributed in-store computing capabilities. SDI can provide resource management for the endpoints so that they can function optimally under various network conditions. SDI and SDN also enable centralized orchestration of data flow from the endpoints and centralized resource management. They provide the capability for flexible composition of network resources at the edge to allow these responsive stores to communicate effectively with corporate headquarters.

CONCLUSION

Intel SDI is critical component required to enable an end-to-end SDDC. It is a reevaluation of system architecture driven by the requirements of business flow, workloads, and specific applications—not by a menu of hardware available to purchase at the moment. SDI has the potential to transform mainstream enterprise datacenters and to displace both current datacenter infrastructure and highly available transaction processing systems by the end of this decade.

The next step to allow cloud datacenters to take advantage of SDI will be the creation of a defined set of use cases and reference architectures where SDI and other “software-defined” technologies can provide significant benefits. MI&S believes that data analytics driven by the Internet of Things (IoT) value chain could be an effective use case for cloud datacenters looking to implement SDDC practices that encompass both SDI and SDN (software-defined networking). The retail inventory optimization and smart factory examples outlined in this paper are two specific usage models that could be enhanced and optimized with SDI.

Network service providers are creating new business models and new services that leverage the power of data analytics from the edge to the cloud. SDI, SDN, and NFV strategies will allow service providers more easily connect the endpoint (sensor) to the backend analytics platform. This will also allow the service providers to morph and adapt quickly to changing needs, repurpose infrastructure with agility, and deploy their control and data plane with flexibility based on the cloud and networking technologies.

Industry leaders should work with leading edge customers in these vertical segments to implement SDDC practices in their IoT value chains that span from the endpoint to the datacenter. Success stories with specific examples of how SDDC practices enhance these environments and create bottom-line value will be critical to help drive broader adoption.

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHORS

Gina Longoria, Analyst at Moor Insights & Strategy
Paul Teich, CTO & Senior Analyst at Moor Insights & Strategy

REVIEW

Patrick Moorhead, Founder, President, & Principal Analyst at Moor Insights & Strategy

EDITOR / DESIGN

Scott McCutcheon, Director of Research at Moor Insights & Strategy

INQUIRIES

Please contact us here if you would like to discuss this report, and Moor Insights & Strategy will promptly respond.

CITATIONS

This note or paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

DISCLOSURES

Intel is a research client of Moor Insights & Strategy, and this paper was commissioned by Intel. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.