Accelerating Cloud-Native Deployments and Economic Returns

Using Dell Telecom Infrastructure Blocks as the Foundation of CSPs' Distributed Clouds

Paul Parker-Johnson

Sponsored by

Dell Technologies

The Industry's Goal: Realizing the Promise of Cloud-Native Designs in CSPs' Network Deployments

Cloud-native architectures create numerous, seemingly limitless benefits on many levels for their adopters and with good reason: they have been the cornerstone of generating remarkable returns for the world's hyperscale cloud providers, their partners, and their customers for over a decade.

Communications service providers (CSPs) worldwide have understood they too can tap into the powerful benefits of cloud-native designs by adopting them in their networks and application deployments. By harnessing openness and modularity of cloud-native designs, they can expand the variety and the appeal of their services and extend their reach to new partners, customer groups, and geographic areas. They can accelerate the pace of their innovations by adopting new technologies faster. They can take advantage of improved economics when developments simplify operations or reduce the cost of an implementation. They can increase their competitive advantage and expand their value to their stakeholder groups.

The Challenge: Achieving the Benefits of Cloud-Native Architectures while also Realizing Efficiencies in Design, Integration Operation

It is **true** that open architectures provide flexibility in solving unique problems and creating new opportunities in both the shorter and longer terms, **but** they also bring with them a challenge of integrating the elements of a design into a solution that is efficient to deploy **and** operate at scale and retain the flexibility to adapt efficiently over time to new applications and technical advances.

In CSPs' distributed infrastructures, this challenge is pronounced. The scope of their footprints is immense, and the number of deployment environments into which they must integrate their offerings is remarkably diverse. With their evolution to a cloud-native infrastructure (as used in 5G) and cloud-native implementation of their services (as in enhanced IoT, consumer and enterprise applications), they are moving irreversibly down a path that requires the versatility of the cloud in the platforms they deploy.

The conundrum they face in obtaining this flexible foundation is that the variety of combinations in hardware, software, integration, and deployment they require is great enough that they cannot afford to evaluate this range of their options effectively with their own personnel (as talented and numerous as they may be). Their ability to create the integrations they need on their own is constrained. They often cannot proceed as rapidly or as extensively as competitive circumstances and customers' requirements demand.

CSPs would be ecstatic to have a way to eliminate or minimize the costs and the complexity with which they are wrestling when considering this number of options.

The Answer: Pre-Engineered, Integrated Combinations of Cloud Platform Infrastructures Focused on the Use Cases and Configurations CSPs Most Commonly Support

CSPs' infrastructures in the future will contain a number of **cloud-native technology stacks** distributed to locations across their footprint. These stacks are composed of layers of technology that each provide functions essential to the CSPs' operation. These include modular cloud-native **infrastructure** layers providing a general-purpose platform for delivering their services, as well as **unique software workloads** at the higher layers of the stack in line to support a particular offering.

On the whole, the infrastructure layer is intended to provide general-purpose capabilities needed by higher level workloads[1]. Although this lends itself toward a mix and match, plug and play philosophy, there remains a variety of alternatives available to CSPs to provide the hardware and software infrastructures they require. For example, hardware and software capabilities required in a remote edge node differ from those required in a centralized national or regional operations data center.

[1] Foundation layers in CSPs' technology stacks generally support the reference architectures found in earlier NFV and more recent cloud-native technology frameworks. These fundamentally rely on an underlying set of general- purpose hardware for compute, storage, and networking; and a set of virtualized, software-based equivalents running on that hardware to support virtualized, cloud- native instances of network functions (in either virtual machine or container-based modes); and a supporting set of management and orchestration functions. It is not the purpose of this paper to explore the details of these alternative reference architectures and what they support; those topics have been amply covered in other work. We assume a basic understanding of those frameworks in this document, as they are all supported by the cloud foundation technologies (both hardware and software) that we are referencing in this work.

Also, capabilities available in alternative suppliers' cloud management software differ in completeness, deployment mode, and operation. Even though capabilities in infrastructure layers are designed to converge on a set of general-purpose functions, evaluating the best combination to deploy in a given CSP requires looking at a number of permutations that make evaluations in-house in a CSP's labs either costly and time-consuming or in some other cases simply infeasible.

One approach to resolving this conundrum is for the CSP to acquire a set of use case-focused, engineered, prevalidated, preintegrated cloud infrastructure implementations designed to meet the requirements of the most common use cases they support. Such combinations can be inserted as preassembled modules into configurations they are deploying. This is an example of the development, integration, and deployment approach referred to as **shift left** in the software development world, in this case applied to CSPs' cloud platform infrastructures. Shift left refers to moving the integration, testing, and validation work related to a new solution being developed *back* in the solution delivery workflow toward its developers, causing integration and readiness testing of the module to be done earlier in the cycle than would have been the case in more conventional, separated development and test. In this mode the time to prepare and validate a new solution for deployment is reduced, and the time to results is accelerated. The goal of applying this model to the foundation layers of CSPs' cloud technology stack is to save substantial amounts of time and money in preparing the stack of its choice for deployment.

Enter **Dell Telecom Infrastructure Blocks** as the foundation layer in CSPs' distributed clouds, delivering the flexibility and efficiency needed at each life-cycle stage. In working with its CSP customers and cloud software partners, Dell has wrestled with this cloud infrastructure testing, deployment, and life-cycle management challenge directly and has created an innovative approach to addressing it, retaining the benefits of cloud-native architectures from start to finish, top to bottom, and end to end.

Building on its Telecom Multi-Cloud Foundations framework[2], its supply-chain operation, and its customer support and professional services offerings, Dell and its cloud software partners are developing a portfolio of use case-focused, preengineered, preintegrated, prevalidated solutions focused on eliminating the downsides of cloud platform integration and life-cycle management, at the same time contributing to CSPs' accomplishing their efficiency, agility, time to market, and innovation goals.

[2] https://www.dell.com/en-us/dt/industry/telecom/telecom-foundations.htm#accordion0

Dell has named the new offerings **Dell Telecom Infrastructure Blocks**. The concept is, when a CSP has decided which use case and infrastructure stack it plans to support and has identified the size and other characteristics of the deployment (in conjunction with Dell and Dell's partner sales and support) it orders the Infrastructure Block appropriate for the use case as a ready-to-deploy bundle from Dell. It can proceed to other aspects of integration into broader network and application delivery systems without having to work on validating the cloud foundation hardware and software with its own staff, funding, and labs. The Infrastructure Block is delivered ready to go.

In this way, design, integration, and validation of the CSP's cloud infrastructure stacks has shifted left and been accelerated by the work that Dell and its partners have performed. The CSP's time to deploy into new serving areas is reduced, and its opportunity to introduce new technology and new network and application offerings to its customer base is accelerated.

To make this less abstract and illustrate how results from using Telecom Infrastructure Blocks are achieved, we describe Dell's approach to creating and delivering them in more detail. We outline the general approach Dell and its software partners are using to deliver Infrastructure Blocks and then profile the initial offering being brought to market using the approach: Dell Telecom Infrastructure Blocks for Wind River, a preintegrated suite of Infrastructure Blocks using market-leading implementations in both hardware and software to accelerate deployment of cloud-native infrastructures in CSPs' 5G RANs.

Implementation of the Infrastructure Blocks Concept

The concept of Infrastructure Blocks is shown in the diagram of a 5G infrastructure in Figure 1. The general case of the design is depicted in the rectangle at lower left. The solution contains Dell's hardware, life-cycle management software for hardware infrastructure deployments (Dell's Bare Metal Orchestrator), and modules from its software partner Wind River's Studio cloud software infrastructure suite, supporting 5G. The use case orientation of Infrastructure Blocks is shown by their alignment with 5G RAN, edge, and core sites just above the Infrastructure Blocks rectangle, indicating the use of Infrastructure Blocks at any point in that continuum. For these cases, Dell and its partners use their in-depth knowledge of 5G from deep industry and customer engagement, and their products' alignment with the use cases' requirements to right size Infrastructure Blocks to meet the needs of a deployment.

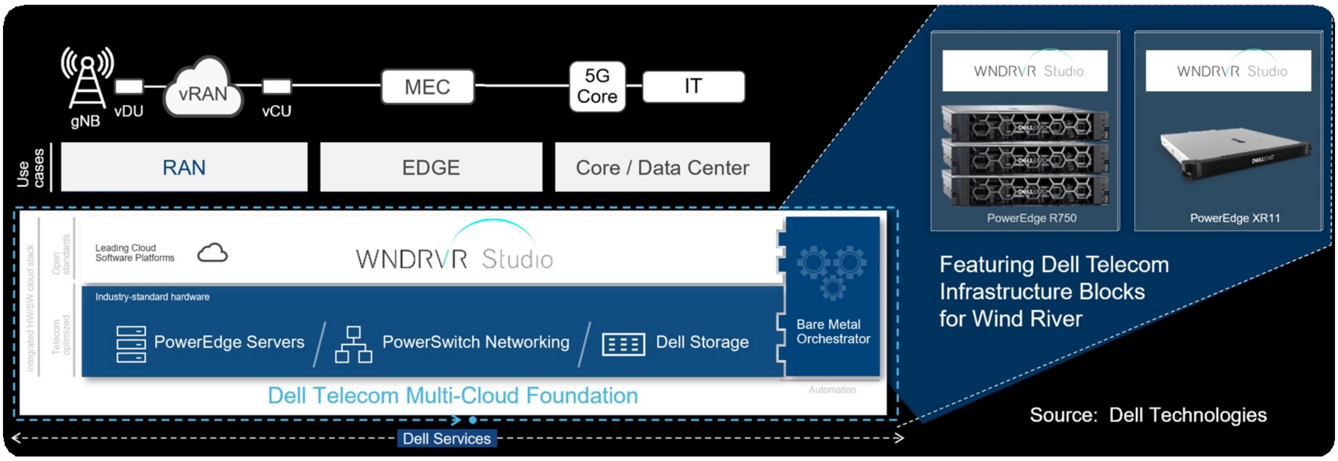

Figure 1. Dell Telecom Multi-Cloud Foundation with Infrastructure Blocks for Wind River

The general approach Dell uses to **design, validate, and deliver** Infrastructure Blocks is illustrated in Figure 2. The work is shown flowing generally left to right. Design, validation, and assembly of the Infrastructure Blocks is performed in Dell's **Solution Engineering Labs** (on the left). When a block is validated and ready to deliver, the package is provided to Dell's supply chain operation, ready to deliver via customer orders to the CSP's labs (on the right). Adaptation to support unique requirements in a given customer's deployment is supported by Dell's investment in a series of **Open Telecom Labs** (OTEL Labs) shown as a branch in the delivery flow of the Infrastructure Blocks into Dell's supply-chain operations. Working with customers and partners in the OTEL context allows fine-tuning of the Infrastructure Block to a given CSP's needs, marrying the efficiency of Infrastructure Block development with the nuances present in meeting a CSP's specific requirements. Continuous improvement in the design of Infrastructure Blocks is enabled by ongoing feedback to the process throughout the life of the solution.

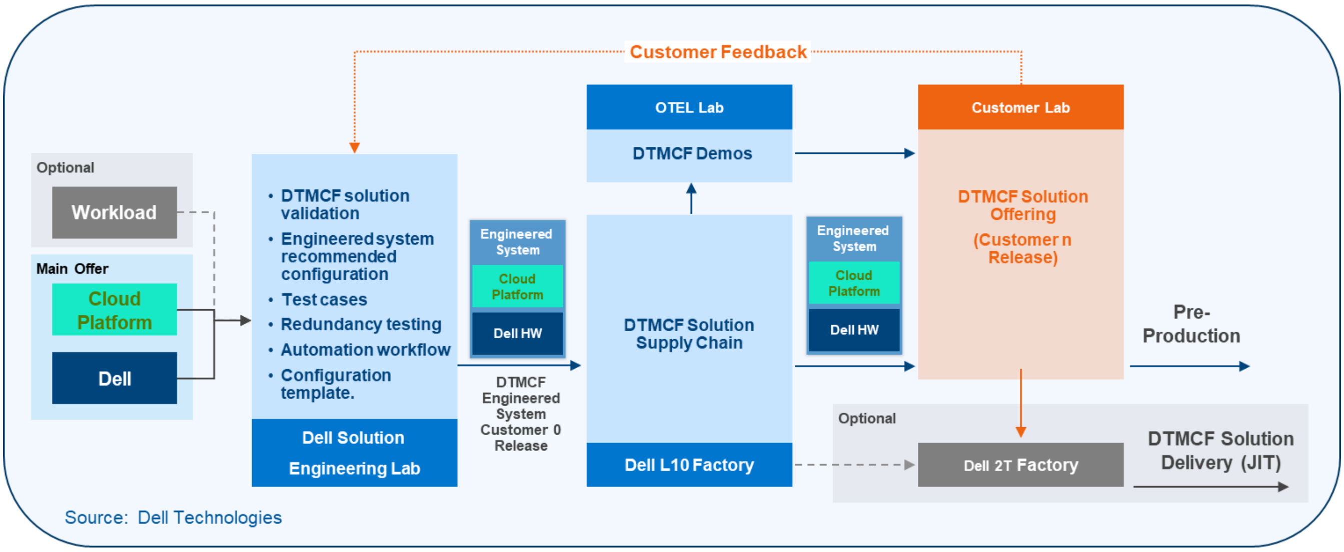

Figure 2. Design, Test, Validation, and Delivery of Infrastructure Blocks for CSPs using Dell's Global Solution Delivery Operation

Having shifted the design, development, and integration of the stacks to the left, Infrastructure Blocks continue to create benefits throughout the life cycle of their deployment (Figure 3).

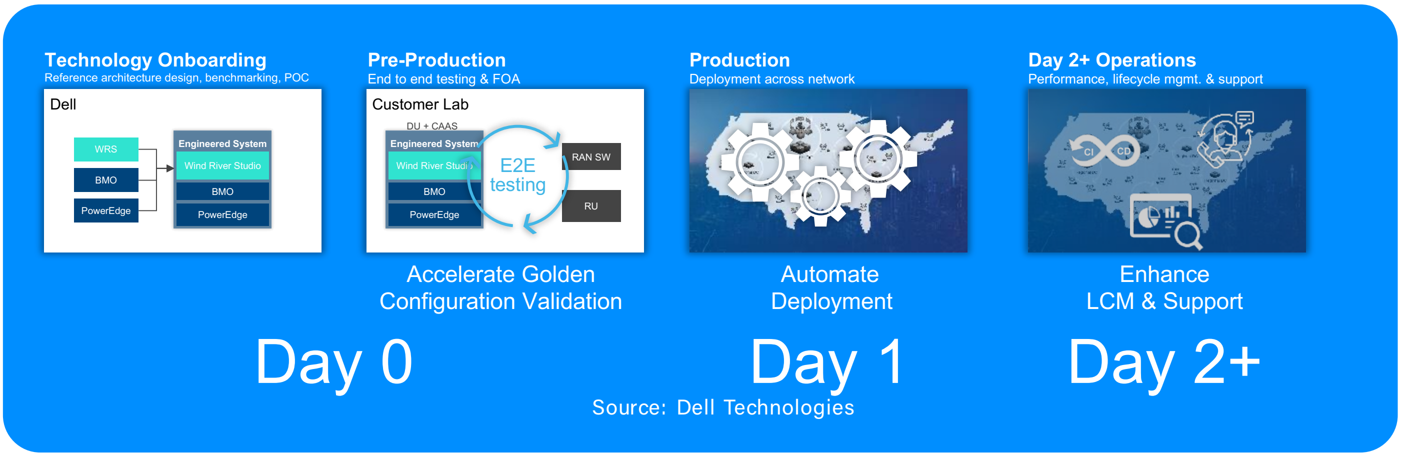

Figure 3. Use of Infrastructure Blocks to Simplify Operations, Create Flexibility, and Accelerate CSPs' Ability to Innovate

We have described Day 0 benefits provided by Infrastructure Block development (shifting design and validation work for the CSP's environment left into Dell's Solution Engineering Labs, preintegrating packages in the labs, and refining designs for the CSP's requirements in OTEL labs). An additional Day 0 benefit comes from Dell and its partners using use case-oriented, field-proven **design guidelines** for the use cases being focused on. These guidelines strengthen the fit of Infrastructure Blocks with the cases and accelerate the preparation of the Infrastructure Blocks for deployment so they are fit for purpose and right-sized from the start.

Downstream from Dell's Solution Engineering Labs, advantages in Day 0 work continue into the CSP's test and integration lab in the preproduction stage. Here, Infrastructure Blocks are integrated with other technology elements of the CSP's deployment, including radios, virtualized network functions, adjacent network domains, CSPs' OSS tool sets, into configurations ready to deploy into production. Infrastructure Blocks accelerate this work by having the cloud infrastructure stack already validated in Dell's solution labs.

Continuing the flow, **Day 1 deployment** into the CSP's footprint is dramatically accelerated using automation embedded in the solution. This includes automation supporting deployment of Dell's hardware (using Dell's Bare Metal Orchestrator in the Telecom Multi-Cloud Foundation) and automation in the cloud software platform being used (such as the Wind River Studio Conductor and Cloud Platform components in Wind River's Studio suite). Tasks are simplified at every step and in every site, including parameters and configurations used at a global level for managing the entire infrastructure and those stored locally to support the requirements of each site.

Day 2+ operations continue to be simplified by the automation included in the Telecom Multi-Cloud Foundation in many ways. First, automation of the complete cloud infrastructure's operation using tools such as Dell's Bare Metal Orchestrator and Wind River's Studio applications (Conductor and Cloud Platform) dramatically streamlines **daily fault, configuration, performance, and security management** in the fully distributed cloud (see the evidence of this in the highlights of our economic analysis in this paper). Second, in alignment with cloud-native design principles, automation in the Infrastructure Blocks supports **integration of the Infrastructure Blocks into a CSP's CI/CD pipeline** using professional services. This significantly streamlines configuration management, contributes to insight-driven, intent-based operations using analytics and automation loops closed to varying extents and allows the CSP to be dramatically more responsive to changes made in its network and application offerings.

A third important benefit of Infrastructure Blocks is significant simplification of technical support. Dell and its partners have implemented a **one-call support model** for managing technical support for Infrastructure Blocks. Whether a question relates to Dell or its partner's component, CSP teams can access the support process using a single point of contact managed by Dell. Tiers of technical support behind that can be accessed and brought to bear on the customer's behalf as required. In this way, CSP teams are spared the overhead that can be involved in coordinating support responses between the companies whose products are deployed. A faster path to resolutions is achieved.

Illustrating the Benefits of Using Infrastructure Blocks in a CSP 5G RAN

As shown in Figure 1, Infrastructure Blocks are well aligned for use cases in 5G, from core to RAN and edge. Over time Dell and its software partners[3] will supply a full portfolio of engineered, preintegrated Infrastructure Blocks across that footprint.

For the first Infrastructure Blocks deliverable, Dell and its partner Wind River focused on a critical domain in 5G operations: cloud-native infrastructures for the data, control and management plane of the 5G RAN in varying Distributed Unit (DU) and Central Unit (CU) configurations. The offering is called **Dell Telecom Infrastructure Blocks for Wind River.**

Supporting RANs in 5G with virtualized, cloud-native operations is one of the keys to helping 5G realize its potential. The number of options available in RAN design for 5G makes it imperative that infrastructures supporting them are flexible, pervasively automated for distributed operation, and maximally efficient in terms of the resources they consume and the performance they deliver.

Both Dell and Wind River have applied themselves extensively to addressing the functional requirements of a successful 5G RAN. Dell has put its hardware expertise to work in designing its **PowerEdge XR11**, a creative combination of general- purpose computing with standards-based support of 5G in a small server footprint, which is incorporated in the Infrastructure Block.

By working with Wind River, Dell has chosen a market leader with deep experience in real-time applications and a wide variety of cloud infrastructure offerings ranging from larger-sized, core network implementations to very efficient, small footprint distributed infrastructures. Its offerings are consistently hardened for reliability, performance, and highly automated operations. Its software is open for integration with a wide range of network functions and a rich array of management and orchestration designs. Wind River has been at the forefront of RAN and edge designs with the **StarlingX** open-source edge software platform[4]. It has also been a leader in scalable cloud software infrastructures for CSPs, starting with Titanium Cloud, evolving into its leading edge, cloud-native

[3] Including Wind River, Red Hat and others.

[4] https://www.windriver.com/studio/operator/starlingx

Wind River Studio portfolio of cloud software infrastructures. Wind River has proven its readiness for use in 5G RAN and edge deployments with its selection by major Tier 1 CSPs for use in their 5G RAN and edge operations in North America and Europe. The combination of Wind River's Studio and Dell's hardware portfolio is a leading edge, feature-rich combination for the initial Infrastructure Blocks solution.

Figure 4 illustrates the modularity of Infrastructure Blocks for options available to support differently sized 5G RANs, small, medium, and large. Blocks are configured and integrated using different amounts of memory, storage, and CPU (among other components) in the hardware and differing compositions of the Wind River Studio software, depending on placement and functionality in the deployment. The size of an installation can scale up, if needed, by adding instances of the Infrastructure Blocks, controlled and managed by the same orchestration software.

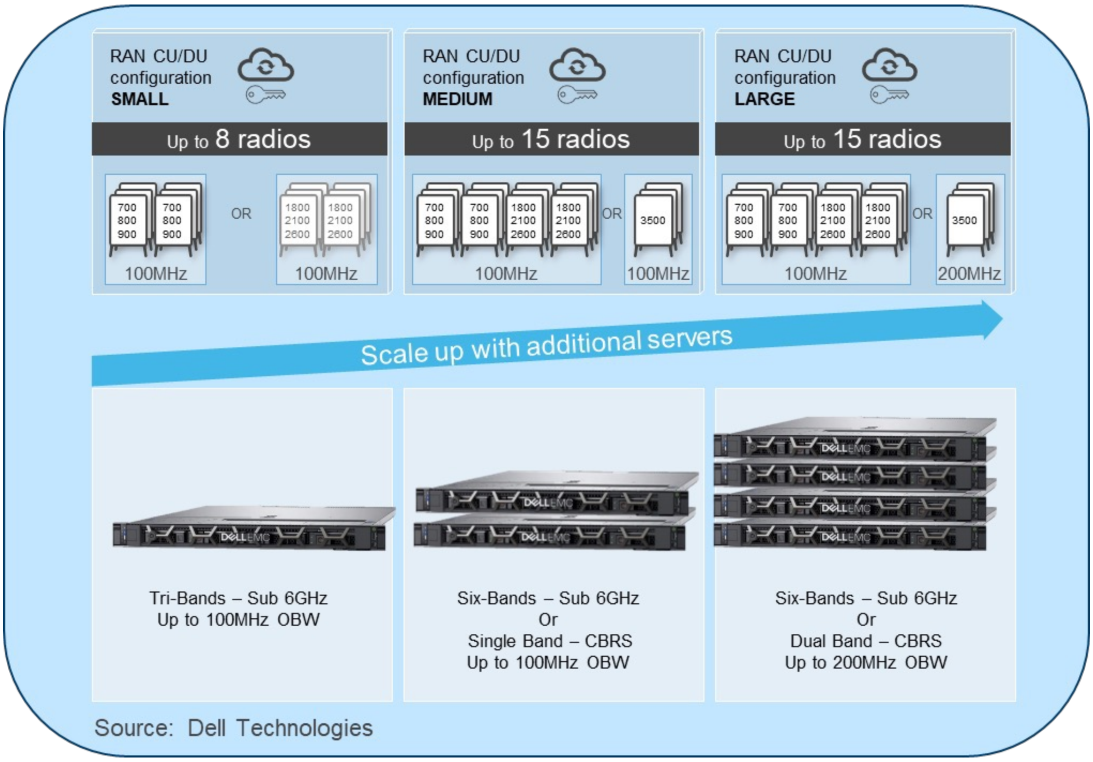

Figure 4. Modularity of Infrastructure Blocks in 5G RAN Configurations

As we previously described, the benefits achievable using Infrastructure Blocks are amplified when combining this fundamental modularity with the efficiencies created by the automation built into them at each layer of the stack.

The Economic Impact of the Model

To calibrate the extent of these benefits ACG Research performed a **life-cycle analysis of the economics** inherent in the design, deployment, and operation of cloud-native infrastructure in a CSP's 5G RAN in two distinct modes of operation. The first mode performs the life cycle of the tasks using separately supplied hardware and software offerings, integrated, validated, deployed, and operated as separate layers by the CSP. We have referred to this as deployment *without* Infrastructure Blocks. The second performs all of the tasks in the life cycle using Telecom Infrastructure Blocks for Wind River; we refer to this as deployment *with* Infrastructure Blocks.

The analysis considers the deployment of a 5G vRAN in a network representative of a European Tier 1 operator's infrastructure or a comparably sized region in a North American operator's footprint[5]. Cell sites are assumed to be deployed in a variety of locations, including dense urban, urban, suburban, and rural. A mix of radio carriers and antennae suited to the needs of each site type is assumed. In this analysis CUs and DUs are centralized in edge data centers for dense urban, urban and suburban sites, and for rural sites we assume DUs are located at the cell site with corresponding CUs installed in edge data centers. The deployment is analyzed over the course of five years. The number of sites begins in year one with 18,000 and concludes in year five with 31,000.

The number of servers deployed to support a given site (and their size) varies with the environment they are supporting (dense urban, rural, etc.). We include Dell PowerEdge XR11s to support DU operations and Dell PowerEdge R750s for CUs. Components of the Wind River Studio software are included based on the function of the node. All of these are standard components of the Telecom Infrastructure Block for Wind River. Based on these assumptions, the server count in year one of the deployment is 32,772 and in year five is 55,960.

The analysis compares operational expenses throughout the life cycle of the deployments in each mode of operation (capital expenses in each scenario are the same and thus not worth comparing). The results show **significant savings** operating *with* Infrastructure Blocks versus operating *without* them). Total operations expense savings over five years in this deployment are shown in Table 1.

---

[5] For a more in-depth description of our analysis, the model on which it is built and the results, reference the companion paper to this research brief, *The Economic Benefits of Dell Telecom Infrastructure Blocks in a vRAN Network Deployment*, December 2022.

	With IBs	Without IBs	Savings Using IBs	RoI Using IBs
Cumulative OpEx after 5 Years	$107 mln	$163 mln	$55.8 mln/34%	109%
Cumulative OpEx after 1 Years	$31.1 mln	$41.5 mln	$10.4 mln/25%	57%

Table 1. Total OpEx Benefits Using Dell Telecom Infrastructure Blocks for Wind River in a 3,000 Cell Site CSP 5G RAN after 1 and 5 Years

More detail on individual areas of cost and saving is contained in Figure 5. This chart is exported from the model developed for this scenario using ACG's Business Analytics Engine. The chart compares expenses incurred throughout the deployment's life cycle in operations *with* Infrastructure Blocks (blue) and *without* Infrastructure Blocks (red).

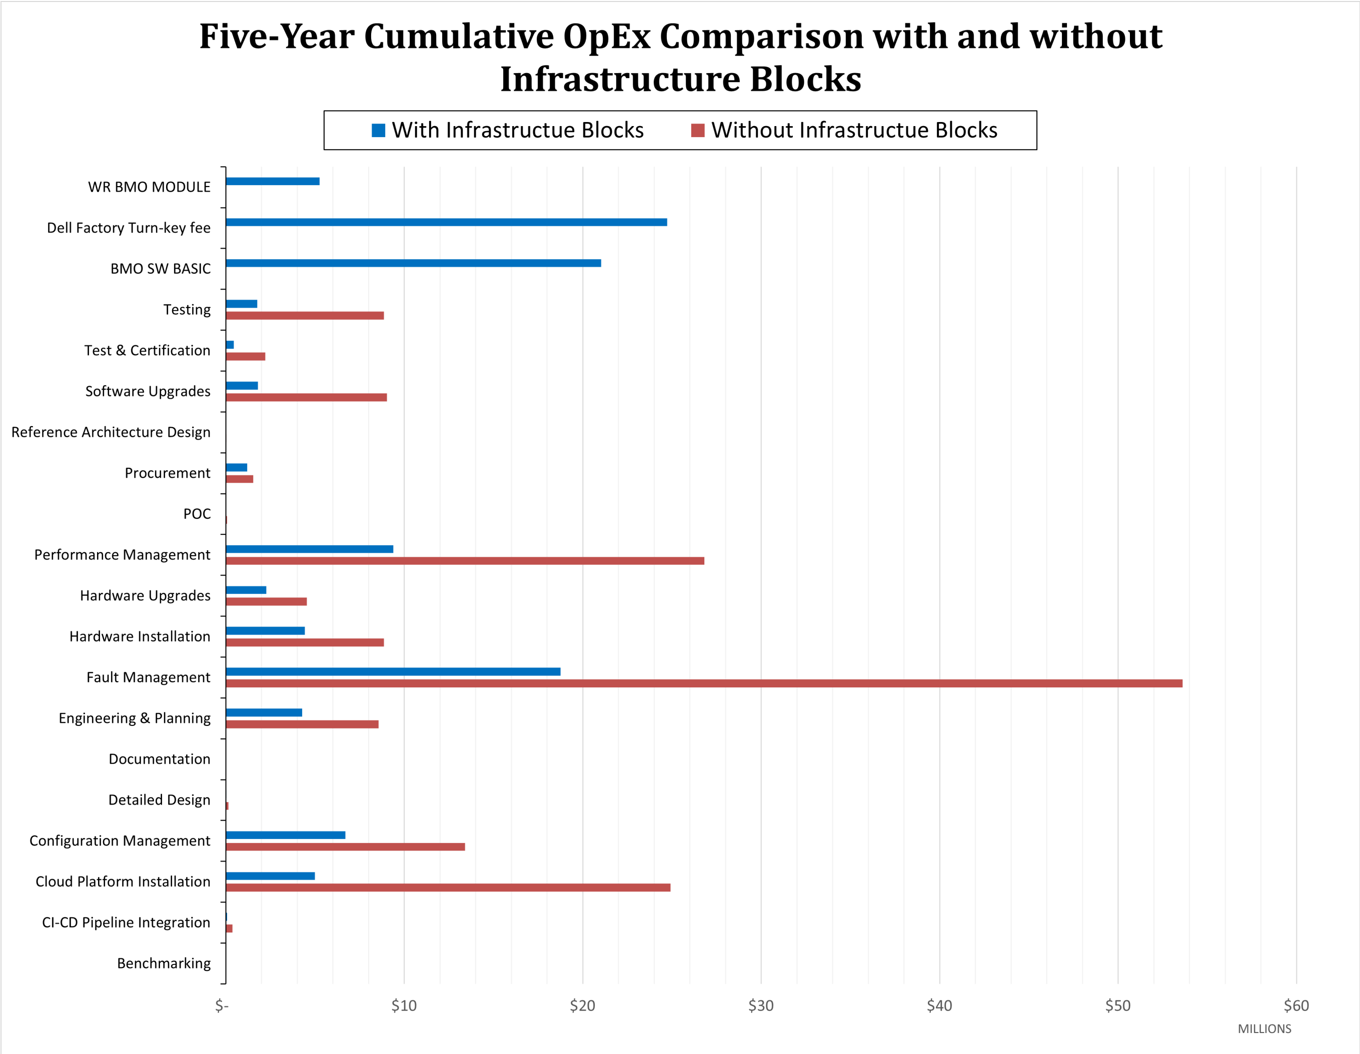

Figure 5. Additional Detail on the Sources of Savings in Operational Expenses Using Infrastructure Blocks

Three categories of expense stand out as significantly more costly in deployments without Infrastructure Blocks: fault management, performance management and cloud platform installation. The first two are the largest because of the frequency of their occurrence during ongoing Day 2+ operations and the challenges introduced for CSPs when coordination between the vendors providing the infrastructure to them is fully the CSP's own responsibility (versus in the case *with* Infrastructure Blocks where Dell has provided the integrated platform to the CSP and has eliminated the coordination responsibility by providing an integrated platform to start with and a single point of contact for support in ongoing operations). The third largest difference (in cloud platform installation) arises because Dell has preintegrated and prevalidated the Infrastructure Blocks versus the case without Infrastructure Blocks where the CSP is doing the integration and the installation all on its own.

There are numerous other categories of expense in which the relative level of cost experienced with Infrastructure Blocks is significantly lower than the costs experienced without them, though the total dollar amounts of the savings are smaller than in the top three categories we identified[6].

Two other categories of cost that involve expenses incurred *with* Infrastructure Blocks that are not incurred in the scenario *without* are worth explaining to help in clarifying the whole context. These are the Dell Bare Metal Orchestrator and the factory turn-key fee that are shown in blue at the top of the chart. These are each expenses that a CSP pays for the Infrastructure Blocks that are not incurred when the CSP is not using Infrastructure Blocks. BMO software is used as a fundamental component of the Infrastructure Blocks that helps simplify and automate operations overall, contributing significantly to savings. The factory turn-key fee is a charge folded in to the Infrastructure Blocks' pricing that represents Dell's costs in producing the integrated solution.

In total across all the categories of expense the differences amount to the 34% lower cumulative operating expenses and 109% RoI over five years when using Infrastructure Blocks, which are highlighted in Table 1.

[6] Refer to our companion Economic Analysis report (Footnote 5) for more detail on each of these.

Creating Increased Value over Time with Telecom Infrastructure Blocks Engineered by Dell and Its CSP Ecosystem Partners

The economic analysis clearly demonstrates the significant impact of using Telecom Infrastructure Blocks in a CSP's 5G RAN. Understanding the value of using cloud-native designs as the foundation of CSPs' operations moving forward, Dell and its partners have also wrestled in depth with the challenges of retaining the openness and flexibility of cloud-native designs while also realizing improvements in the efficiency of their integration, deployment, and support over time. Proactively, they have applied their combined insights into creating an improved path forward for resolving this challenge with their Infrastructure Blocks initiative.

Infrastructure Blocks provide CSPs with the best of both worlds. They retain the openness and modularity of cloud-native designs in a wide range of the most common CSP use cases. And they achieve significant, sustainable improvements in the economics of designing, integrating, deploying, and operating CSP distributed clouds based on expertly engineered, preintegrated, prevalidated, and deployable infrastructure blocks that are ready to use in CSPs' distributed clouds.

In our analysis of Dell Telecom Infrastructure Blocks for Wind River in a representative CSP 5G RAN implementation we found that CSPs can achieve a 34% savings in the operational expenses they incur over the five-year life cycle of their deployment. The RoI of using Infrastructure Blocks in the 5G RAN case is 109% after the full five-year cycle.

These findings bode well for the advantages Infrastructure Blocks can provide to CSPs as the number of use cases included in the portfolio expands to 5G core, edge, and other prominent CSP implementations. Dell and its partners have made a meaningful, distinct, and sustainable contribution to the state of the industry's practice in putting cloud-native architectures to work efficiently and at scale in a wide range of CSPs' deployments.

Paul Parker-Johnson

Paul Parker-Johnson is ACG's research lead in cloud and virtual system infrastructures and their use in private, hybrid, edge and multi-cloud designs. His work encompasses physical, virtual and cloud-native infrastructures; use of intent-based modeling and automation; fusion of OT and IT in distributed cloud designs; and support of use cases in 5G, industry 4.0, and digital transformation.