



Scaling Productivity and Innovation on the Path to Exascale with a “Team of Teams” Approach

Elaine M. Raybourn¹(✉), J. David Moulton², and Aimee Hungerford²

¹ Sandia National Laboratories, Albuquerque, NM 87185, USA
emraybo@sandia.gov

² Los Alamos National Laboratories, Los Alamos, NM 87545, USA
{moulton, aimee}@lanl.gov

Abstract. One of the core missions of the Department of Energy (DOE) is to move beyond current high performance computing (HPC) capabilities toward a capable exascale computing ecosystem that accelerates scientific discovery and addresses critical challenges in energy and national security. The very nature of this mission has drawn a wide range of talented and successful scientists to work together in new ways to push beyond the status-quo toward this goal. For many scientists, their past success was achieved through efficient and agile collaboration within small trusted teams that rapidly innovate, prototype, and deliver. Thus, a key challenge for the ECP (Exascale Computing Project) is to scale this efficiency and innovation from small teams to aggregate teams of teams. While scaling agile collaboration from small teams to teams of teams may seem like a trivial transition, the path to exascale introduces significant uncertainty in HPC scientific software development for future modeling and simulation, and can cause unforeseen disruptions or inefficiencies that impede organizational productivity and innovation critical to achieving an integrated exascale vision. This paper identifies key challenges in scaling to a team of teams approach and recommends strategies for addressing them. The scientific community will take away lessons learned and recommended best practices from examples for enhancing productivity and innovation at scale for immediate use in modeling and simulation software engineering projects and programs.

Keywords: High performance computing (HPC) · Exascale · Team of teams · Computational science & engineering (CSE) · VUCA · Organizational productivity

1 Introduction

Since the collapse of the Soviet Union in the early 90s, the term VUCA (volatile, uncertain, complex and ambiguous) has been used to describe current and future contexts for militaries, industry, and education. However, it is rarely acknowledged that modern *science* driving innovation in each of these domains is *also* VUCA. It is the *scientific discovery*, often executed in rapidly changing and shifting VUCA contexts, that drives advancements in technology and processes that in turn create the

The original version of this chapter was revised: The copyright holder name has been corrected. The correction to this chapter is available at https://doi.org/10.1007/978-3-030-22338-0_39

opportunities for breakthrough solutions in computational science & engineering (CSE), modeling, and simulation. Science *is* VUCA, and the teams whose scientific discoveries must anticipate 21st Century challenges are also operating in environments that relentlessly shift in tempo, complexity, and scale.

However, while living in an increasingly VUCA environment is certain, it is also evident that *not everything is unpredictable* [1]. When it comes to drivers for change, innovation, and subsequent modernization the U.S. whole of government must be poised to prepare for, and address stressors to national security. The U.S. Department of Energy (DOE) and the DoD—are, of all the U.S. Federal Executive Departments, in synergistic positions to apply advancements in science and exascale computing to modeling and simulation for national security impact on global trends such as urban concentration, demographic shift, climate change, and technological surprise. Together, federal agencies can prepare for these known global trends that will, and have already begun, to shape our near and distant futures.

1.1 What Is Exascale, and Why Is It a National Security Grand Challenge?

If high performance computing (HPC) today functions above a petaflop (computing speed equal to one million million (10^{15}) floating-point operations per second) then exactly what is exascale computing? Exascale computing is a new class of high-performance computing systems capable of 10^{18} floating point operations per second [2] and whose power is measured in *exaflops*, or computing speed equal to one billion billion calculations per second, and one thousand times more powerful than today's petaflop machines. By 2023, exascale computing will enable us to solve incredibly complex research problems, such as simulating the human brain, hypersonic aircraft, and nuclear tests; predicting space weather, cracking encryption codes, advancing medical research, analyzing 22 million Veteran health records, and modeling the Earth's climate and the global economy [3–5]. Put in layman's terms by the Secretary of Energy, “the jump to exascale would be like going from a flip phone to the latest smartphone or from dial-up to 4G internet speed” [3]. Nevertheless, in the midst of much enthusiasm, there are several pressing trends contributing to the recent sense of urgency. First, impacts on society, energy security, and our everyday lives contribute to the collective pressure to succeed. Second, data is growing exponentially, and while increasingly accessible, is overwhelming. Third, computational power is growing (and largely compounding the challenge). And fourth, flexible, extensible and agile interdisciplinary approaches and solutions seem most promising when they move toward a capable exascale computing ecosystem.

2 The Path to Exascale

The High Performance Computing Modernization Program (HPCMP) is not a well-known program although it was established in 1992 to modernize the HPC capability of DoD laboratories [6]. The HPCMP was “*assembled out of a collection of small high performance computing departments*, each with a rich history of supercomputing

experience that had independently evolved within the Army, Air Force, and Navy laboratories and test centers” (U.S. Army Corps of Engineers HPCMP fact sheet, italics ours). The program’s capabilities have been used to accelerate acquisition decision-making, simulate wind tunnels, decrease development timelines [6] for the design of helicopters, and for analyzing weather patterns [7].

As the lead federal agency for fundamental research in energy, and the “single largest supporter of basic research in the physical sciences in the United States” (<https://science.energy.gov/about/>), the DOE Office of Science has a core mission is to move beyond current HPC capabilities, such as those employed by the HPCMP, toward a capable *exascale computing ecosystem* that accelerates scientific discovery and addresses critical challenges in energy and national security. An *ecosystem* is defined as a group of independent but interrelated elements comprising a unified whole, and the same definition is applied to an *exascale computing ecosystem* in the present paper. The capable exascale computing ecosystem being developed by the DOE Exascale Computing Project (ECP) is an exascale ecosystem which includes applications, system software, hardware technologies and architectures, and critical workforce development—independent, but interrelated elements that comprise a whole. Finally, the DOE Office of Science maintains a research portfolio that includes advanced scientific computing, basic energy science, biological and environmental research, fusion energy science, high energy physics, and nuclear physics. Through scientific discovery and execution of the mission with its semi-autonomous National Nuclear Security Administration (NNSA), the DOE enables application of nuclear science to modeling, simulation, and testing at exascale.

The DOE and modeling, simulation, and test communities can employ lessons learned from organizations leveraging interdisciplinary teams to address VUCA experiences. Recognizing that the path to exascale will draw upon the talents of a workforce enabled primarily by small teams, we take pages from General Stanley McChrystal’s playbook [8] and apply them to scaling from small teams to a *team of teams*.

2.1 VUCA Requires Organizational Agility that Scales

High performance computing as an international community of agencies, national laboratories, academia, and industry has largely been defined by the ad hoc assembly of small, agile computational science & engineering (CSE) teams [9]. For many CSE teams, past success may have been achieved by working shoulder-to-shoulder on small, trusted teams (e.g. professor, collaborating students and postdocs, or a modest sized research group) that rapidly innovated, prototyped, and delivered. These aggregate teams are often distributed across organizations, physical locations, time zones, language communities, and cultures. As such, it can be a challenge to foster successful collaboration in larger, distributed teams. The VUCA nature of this socio-technical path to exascale introduces significant uncertainty in HPC scientific software development for future modeling and simulation, and can cause unforeseen disruptions or inefficiencies that impede organizational productivity and innovation critical to achieving an integrated exascale vision.

VUCA contexts require networked configurations—those able to shift in execution and optempo, thereby enabling organizational agility and resiliency at scale (see Fig. 1). A networked configuration should facilitate interdependency, cooperation, and coordination at the level of individuals, teams, and organizations. Therefore, we must reconsider hierarchical “command structures” that are no longer effective when faced with asymmetric threats. Hierarchical “command structures” are designed to be efficient, but efficiency is no longer enough to address a new generation of VUCA threats. Many industrial organizational models in place today were designed to allow organizations to better plan, standardize, and predict output, however, VUCA requires a new approach.




Fig. 1. From few to few, toward many to many.

Figure 1 illustrates three types of organizational structures discussed by McChrystal and others [8]: hierarchy, hybrid structure, and network. The hierarchy illustrates the typical “command structure” found in many military organizations. These organizations are characterized by an emphasis on efficiency which leverages silos in predictable ways. On the other hand, the network structure at the right in Fig. 1 focuses on increased interconnectedness and flat, flexible communication. These organizations are more agile, resilient, and well-suited to uncertain environments. Finally, the hybrid structure in the center of Fig. 1 can be employed by hierarchical “command structure” organizations as they adjust their practices toward a more networked system. The hybrid organization encourages more communication across silos. Like many organizations, ECP may be characterized as sets of small teams functioning within a hierarchical structure. Therefore, ECP requires a systems approach that will stimulate interconnectedness for productive and innovative, computational software engineering. On the path to exascale, ECP members will likely need to understand the entire interconnected system, not just their individual areas of expertise or those of siloed teams. Harnessing the entire capability of the distributed organization requires information sharing to achieve levels of transparency that may be entirely new to the organization. According to McChrystal and others [8], the command of teams may be more flexible in some organizations than others, but to be truly adaptable in a VUCA environment, much more than a hierarchical structure is required.

3 Scaling from Small Teams to a Team of Teams: ASC Ristra

The DOE Advanced Simulation and Computing (ASC) program supports the tools to underpin the use of simulations in assessing the current and future stockpile. Within the ASC program the Advanced Development, Technology and Mitigation sub-program (ADTM) focuses on the development of new, production simulation tools for the Nuclear Security Enterprise. The VUCA context faced by the team is created by disruptive trends in hardware, programming models and tools, and in the increasing complexity and diversity of research questions these tools must address. The Los Alamos National Laboratory (LANL) ASC Ristra project (named for the threaded chile strings found in New Mexico) is supported by the ASC-ADTM sub-program as part of ECP and focuses on addressing technical surprise for HCP through the development of a Flexible Computational Science Interface (FleCSI) that will limit the impact of volatility and uncertainty [10]. Specifically, FleCSI is the framework and foundation for a software ecosystem of interoperable and composable components that will provide flexibility and adaptability to effectively address a broad range of science and analyses. In addition, it will provide an important abstraction layer that shields developers from the variability in computational platforms.

Thus, a key challenge for the ECP (Exascale Computing Project) is to scale the efficiency and innovation delivered by small teams to the aggregate team of teams illustrated in Fig. 2 [8]. While scaling agile collaboration from small teams to teams of teams may seem like a trivial transition, this paper describes how the path to exascale in HPC scientific software development for future modeling and simulation would be enhanced by a team of teams approach. We use ASC Ristra, an existing DOE ECP team, as an exemplar in the present paper that illustrates key elements necessary to scale a large organization to a hybrid structure (see Fig. 1) and eventually toward McChrystal and his co-authors' approach of a team of teams (Fig. 2). Subsequent sections discuss shared consciousness, empowered execution, and leading like a gardener. We conclude by recapping key take-aways for the scientific community.




Fig. 2. Team of Teams.

3.1 Shared Consciousness

This section introduces the idea of developing shared consciousness among and within teams as an important step on the path toward scaling the productivity of small teams to a large networked aggregate team, or a team of teams [8]. In our own projects such as ASC Ristra, we have been exploring ways to apply lessons learned from the book, “Team of Teams: New Rules for Engagement in a Complex World” by retired US Army General Stanley McChrystal and others [8]. The authors describe a team of teams as a network of coordinated teams who are interdependent and share trust as well as an awareness of a common purpose, and goal. This connectivity of common purpose, trust, awareness of each other, and transparent communication is referred to as shared consciousness [8]. We don’t have to be in project management or a formal leadership role to help our teams reach across rigid organizational structures to achieve shared consciousness.

It has been noted that shared consciousness among teams can enhance organizational productivity and facilitate innovation critical to achieving an integrated vision [8]. Whether we are co-located or distributed, post docs, team members, or group leaders, we can all participate in reaching an integrated vision together. Below we summarize three strategies advanced by McChrystal and others [8] toward developing shared consciousness for and among teams that are distributed and/or co-located.

Identify a Single, Aligning Narrative. Before we can have shared consciousness among our distributed teams, we must first foster a shared mental model of why we are doing what we are doing. That is, before focusing on the technical capabilities to connect, we can take time to engender the willingness to connect [11]. One way to accomplish this is to have a single unifying message that speaks to everyone—consider, each individual, on every team, can see themselves as participants in the story as it unfolds [12, 13]. Repeating this aligning narrative as often as needed, perhaps by opening each meeting with the common purpose, will enable team members to feel comfortable articulating it themselves.

A shared mental model will require, at a minimum, a shared vocabulary. Since the subject matter experts that have come together on the ASC Ristra team represent many different fields of physics, computer science, and HPC, and in many cases, sub-sets of contributors who had worked together on previous teams, the ASC Ristra team was built through the inheritance of smaller software development teams. The inherited teams ranged from academic computational research scientists in astrophysics to LANL production software development teams for multi-scale, multi-physics applications. The software development environments of each group were widely varying on several fronts. Technical fields of expertise were spread across applied mathematics, computer science, and physics with each group bringing its own unique vernacular to the mix. Coding languages and styles from the diverse backgrounds were extremely varied, e.g. Fortran, C, C++, Python and Haskell, and came with code styles that reflected a spectrum of software discipline, from single user research codes to production released community codes. With this diverse background, the ASC Ristra team was faced with engendering a shared vocabulary as a critical step in establishing effective communication and an aligning narrative.

The technical vision of ASC Ristra is to support diversity in science applications and compute architecture through the design and implementation of FleCSI and the ecosystem of supported components (Fig. 3). ASC Ristra uses visualizations such as Fig. 3 to help focus shared consciousness by communicating several key aspects of software design goals and project structure. It makes it easier for team members to understand their roles within and across capabilities, and highlights key mechanisms for mitigating the volatility and uncertainty from external drivers. In addition, visualizing the impact of VUCA on the software design goals as well as project structure gives team members awareness of the breadth of components critical to the software ecosystem, and consequently, the diverse and interdisciplinary nature of the team that is needed to develop and support it. This is an important point in that visualizing VUCA on a regular basis creates shared consciousness of why the vision is urgent, the pressures to succeed, and who can help.

A clear technical vision such as the one illustrated in Fig. 3, and recruitment of strong technical teams are necessary but not sufficient. To develop a successful aligning narrative, ASC Ristra added to their technical vision: embracing their team's diversity and valuing *interpersonal relationships*. These important additions were also key elements in the aligning narrative developed by McChrystal and others [8] and clarified by Fussell and Goodyear [11, pg. 58]. However, it's important to note that fostering interpersonal relationships is the ingredient that is *most often overlooked or dismissed as a low priority* that will simply take care of itself. This is simply not true. By strengthening interpersonal relationships upfront, and with high priority, ASC Ristra is much more likely to succeed. In the following section we expand on how ASC Ristra develops and supports relationships by building trust.

Build Trust. Use a common conceptual model (as shown in Fig. 3) and aligning narrative of the entire, interconnected project to build trust among individuals within a team, and across distributed teams. The terms “within” and “across,” may also refer to vertical and horizontal alignment in hierarchical organizations. Many of us who are part of hierarchical organizations may often perceive building trust within vertical alignment to be easiest, although that may not always be the case. Trust may also be built across horizontal organizational alignment with individuals with whom we share common purpose, language, and conceptual models. In the book, “Great at Work,” Hansen [14] describes trust as confidence in individuals to consistently and reliably deliver high quality work that is expected from them. If our teams have not reached this level of trust, Hansen recommends applying “trust boosters” to specific trust problems as needed. For example, if distributed team members across a project do not know each other well (they may be strangers), the teams may want to try some bonding exercises or techniques to share personal information.

Many of the initial challenges encountered by ASC Ristra required a willingness for both cultural change and adoption of new technical approaches as well as “trust boosters” Hansen [14]. The team acknowledged and openly welcomed lessons learned from the past efforts, because they saw these as keys to helping plot a different course, rather than harbingers of inevitable doom. This attitude shift was accomplished by paying close attention to the make-up of the team. Diversity of experience, personality and technical expertise were actively considered, even at the team leadership level.

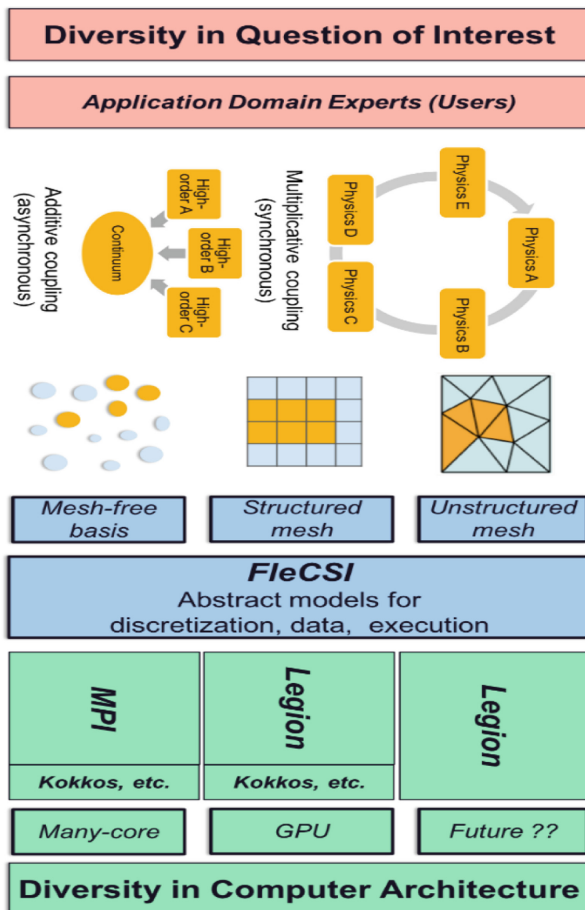


Fig. 3. ASC Ristra schematic of external diversity on the software ecosystem.

A mix of senior and junior staff across a breadth of relevant technical fields was fairly easily achieved. More challenging was finding the right mix of personalities, so that the psychological diversity could accompany that technical diversity.

A “trust booster” the ASC Ristra leadership relies on is *regularly scheduled, informal collaboration time* (in the form of afternoon “Tea” twice a week) to help gauge this balance of diversity, as well as provide a place for creating a shared consciousness within the team. Office co-location had long been a strategy for achieving team cohesiveness at LANL but was not an easy option for ASC Ristra. This was in part due to space limitations, but also because this team required part time engagement from a very large set of subject matter experts. Shared consciousness may be hindered or helped by physical and virtual spaces, and established processes. Creating shared consciousness at scale requires rethinking every process in the organization, including how we approach physical and virtual spaces. The Tea-time concept has been an effective way to overcome lack of co-location, and the Tea-time room has become a

natural gathering place for impromptu sub-team meetings at other times. As a result, the general availability of this collaboration space has been a key ingredient for the team's shared consciousness. The trust booster of "*regularly scheduled, informal collaboration time*" has been adopted by at least one other ECP team known by the authors in the form of afternoon "cake breaks."

Commit to Sharing Information to Develop and Maintain Strong Connections Among Teams. Achieving a team of teams will rely on our ability to foster interdependent teams. Interdependence does not merely rely on co-located, "shoulder-to-shoulder" communication. Interdependence must be built on information sharing that is transparent and efficient, and allowing teams to execute work based on their initiative. Since teams are increasingly distributed and virtual, we can try using computer-supported cooperative work (CSCW) cloud-based development tools and video/whiteboard meeting software environments. The use of information technology will help propagate information across the network and foster more transparent communication [12, 13]. We'll have to be patient—strong connections won't happen overnight.

Developing common tools and standards for communication within and across the interdisciplinary teams comprising ASC Ristra presented a variety of challenges. The staff can be divided into two major categories: those with office space in an open environment with fairly modern technology available to them, and those whose workspace is protected for security reasons with significant restrictions on use of cell phones, tablets and other convenience electronics. This created a significant divide in the comfort level with communication tools. Some staff are more comfortable with face-to-face interaction, while others are insistent about using online communication tools and social media. Common software development environments were also not shared by the teams that came together within ASC Ristra. Several team members were not familiar with using a distributed repository control system like Git and were more accustomed to developing from a tar file exchange with collaborators. In addition, activities with similarity to the stated goals of ASC Ristra have been taken on in the past at LANL with limited success. The team started off surrounded by those who had walked their path before and found only disappointment.

Thus, an important topic in the early informal Tea-time gatherings was the exploration of new CSCW tools and approaches that could effectively support a collaborative environment for this diverse team. Specifically, the ASC Ristra team spent its early months seeking a software development tool stack that could fit its needs and be accessible to all the staff involved. The Atlassian tool chain was ultimately adopted and a series of tutorials targeting its use case within ASC Ristra were offered to the whole team. This type of activity is important for creating a shared consciousness, it requires a significant time investment from each member of the team, and often does not generate "results" in any traditional sense of the term. As such, learning-curve activities like this require significant management support so that staff can be shielded in these early phases from product delivery just for delivery's sake. A by-product of this necessary shielding is the headroom to explore higher risk technical ideas as well. ASC Ristra has found significant benefit from having allowed time to chase several higher risk innovative ideas.

In summary, when scaling from small teams to a team of teams, consider how to bootstrap your efforts with the goal of shared consciousness. While this may sound daunting at first, scaling collaboration found in small teams to an aggregate can be made much smoother within our programs by considering the connectivity of an aligning narrative, trust among interdependent teams, and the use of modern CSCW technology to foster transparent communication, and efficient information sharing. Taking time to build trust and a unified sense of purpose within and across teams are important steps in scaling from a small team to a distributed aggregate, or team of teams.

3.2 Empowered Execution

As we discussed in the previous section, the characteristics of effective, small team are trust, common purpose, shared consciousness, decentralized decisions, and autonomy. *But how can we scale up?*

To scale as a team of teams, organizational configurations should foster the delegation of authority to sub-teams and individuals who can act without explicit senior leadership consultation [8]. Providing opportunities for team members to understand each of the interdependent parts of the system will help create a general awareness that will support each team member's specific area of expertise. Introduce transparency in communication and processes to establish a culture of shared consciousness. The synchronization of team cadence in the hybrid model can also reaffirm the aligning narrative.

The importance of preserving both the autonomy and identity of individuals within the context of a team of teams environment cannot be overstated. For example, the ability to respond quickly to new challenges, the confidence to share information openly in integrated, technologically-mediated forums such as videoconferences, the efficient execution of team or individual-specific activities, and the branding associated with both individuals and team capabilities or components can all be enhanced when approvals from a chain of command are not required. Conversely the unnecessary use of vertical approval chains can severely inhibit or even silence these valuable activities.

As discussed earlier, the challenge is to bring the foundational support implicit in a small team, such as trust, common purpose, shared consciousness to the team of teams. This is critical to an effective team of teams environment as it enables periods of empowered execution between broader synchronization points. In the context of large interdisciplinary software development projects this empowered execution can be thought of as disciplined and focused autonomous development within small team environments that is guided by the aligning narrative and technical vision. Although, it is effectively a small team setting, it may be an ad hoc team that includes liaisons or interdisciplinary sub-teams. In addition, it follows the accepted best practices of the broader team, including, styles, tools, communication protocols, and testing protocols.

ECP projects are actively exploring a range of methodologies that can support this hands-off element of empowered execution, while also supporting the accepted best practices. These approaches include the use of bull pens (open-office arrangements for small teams, possibly over specific timeframes), and virtual scrums. In addition, workflows are evolving that allow teams to tune the balance of supervision and

autonomy, such as the use of pull requests and code reviews in cloud-based source code management systems like GitHub. In this case, significant freedom and creativity can be leveraged early in the process, but the result put forward in the pull request can be reviewed to ensure it meets development goals, accounts for upstream and downstream dependencies, and follows the software standards of a project. In addition, these methodologies and tools still support an “eyes-on” overall situational awareness, as activities are naturally captured by the interaction and logging that development leveraging these tools and workflows creates. Finally, as discussed in the subsequent section, scaling work flows and project processes through empowered execution requires a paradigm shift in the execution of leadership.

3.3 Leading like a Gardener

A successful shift to a team of teams structure that maintains shared consciousness and enables empowered execution must also be accompanied by a fundamental shift in the leadership style. The most useful and intuitive metaphors describing this shift are from chess master to humble gardener [8]. Specifically, as we strive to accelerate scientific discoveries in this VUCA environment, the critical observation is that centralized control is too slow. It hinders cross-fertilization, and stifles innovation. Even with the amazing wealth of real-time digital information available to leaders today, it is unreasonable to expect them to perform as chess masters and effectively use this information to micro manage a reductionist view of teams and resources.

Instead, to create and nurture an environment that enables “smart autonomy” and shared consciousness, leaders must adopt the role of the humble gardener [8]. Specifically, this environment uses the hybrid and network structures to extend and enhance the qualities implicit in the cultures and processes of successful small teams. To be effective this environment must foster trust, transparency, and free-flowing conversations, both within and across teams. When high-quality liaisons are identified and used as critical nodes in cross-team networks, they can play a significant role in supporting this transparent and effective communication. Furthermore, the leader, as the gardener, plays a critical role in sustaining this environment by consistently reaffirming the aligning narrative, continuously demonstrating respect and appreciation for everyone’s contributions, and reinforcing empowered execution with an eyes-on, hands-off, approach [15]. Thus, gardening in reality and in this metaphor, is not passive. The gardener, drives the rhythm and transparency of the team of teams, and builds and maintains the trust needed for cross-functional cooperation.

ASC Ristra started before McChrystal and others wrote “Team of Teams,” and at a time when the uncertainty and volatility of computational science had reached a tipping point and the community was searching for approaches to restore or enhance overall productivity. Lessons learned from previous work on multi-physics applications and frameworks, advances in tools and languages, and the growth of agile methodologies began to fuel the vision of a software ecosystem that would provide the flexibility and adaptability required to meet 21st Century VUCA contexts. During this time, it became apparent that the leadership style needed to change as well. ASC Ristra were among those keenly aware of these changes and adopted many aspects of the humble gardener to support their aligning narrative and effectively develop this new software ecosystem.

4 Conclusion

Take-aways have been provided in each section of the present paper and now we recap key concepts by presenting the following questions and advice for consideration as your own organizations scale with a team of teams approach:

1. What is your aligning narrative? Fussell and Goodyear [11, pg. 61] offer the following question to consider as you create your narrative: *In your organization's current state, how would an aligning narrative be best emphasized and contextualized for your teams?*
2. Do your partnerships need a “trust booster?” Fussell and Goodyear [11, pg. 61] offer the following question as you consider how well your aligning narrative is supported by trust: *Your teams might claim to trust and have open lines of dialogue with one another—but can they demonstrate examples of how this trust is helping accomplish your organization's strategy?*
3. Do you leverage a geographically distributed organization to your advantage? Fussell and Goodyear [11, pg. 108] offer the following question as you consider how well your aligning narrative is supported by trust: *Could your organization leverage both physical and virtual spaces to drive cross-silo coordination and improve communication on strategy?*
4. Do you practice disciplined, empowered execution to scale and adapt your organization's response to emerging threats? McChrystal et al. [8, pg. 291] offer the following advice: *“Individuals and teams closest to the problem, armed with unprecedented levels of insights from across the network, offer the best ability to decide and act decisively.”*
5. Does your organization reflect a model that is out of date? McChrystal et al. [8, pg. 232] offer the following advice: *“Although we intuitively know the world has changed, most leaders reflect a model and leader development process that are sorely out of date. We often demand unrealistic levels of knowledge in leaders and force them into ineffective attempts to micromanage...Persuading teams to network with other teams will always be difficult, but this is a culture that can be planted, and if maintained, can flourish.”*

The present paper described how HPC CSE software development for future modeling and simulation projects would be enhanced by fostering a team of teams environment. Our discussion of the team of teams approach conceived by General Stanley McChrystal and his co-authors highlighted shared consciousness, empowered execution, and leading like a gardener through the use of an exemplar from DOE ECP. The path to exascale introduces significant uncertainty in HPC scientific software development and can cause unforeseen disruptions or inefficiencies that impede organizational productivity and innovation critical to achieving an integrated vision. We hope our discussion has illuminated a key element that while *often under-reported*, is integral to the scientific discovery process—workforce development toward enhanced organizational productivity. According to the ECP vision statement, an exascale ecosystem will incorporate applications, system software, hardware technologies and architectures, along with critical workforce development. The success of

the program depends on its people. Essentially, it is necessary to scale a shared consciousness intrinsic to smaller teams, which enables autonomous periods of heightened productivity, to a productive and sustainable ecosystem of teams of teams.

Acknowledgements. *Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. Images in Figs. 1 and 2 copyright by NTESS LLC and reproduced with permission. **This work was supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research (ASCR), Office of Biological and Environmental Research (BER), and by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. Figure 3 reproduced with permission. Special thanks to the ADTM-ASC Ristra project for their willingness to be highlighted as an exemplar for the approach, challenges, and successes discussed in the present paper.

References

1. McNulty, E.J.: Leading in an increasingly VUCA world. Strategy-business.com, 27 October 2015. <http://www.strategy-business.com/blog/Leading-in-an-Increasingly-VUCA-World?gko=5b7fc>. Accessed 4 Feb 2019
2. Reed, D.A., Dongarra, J.: exascale computing and big data. Commun. ACM **58**(7), 56–68 (2015). <https://cacm.acm.org/magazines/2015/7/188732-exascale-computing-and-big-data/abstract>. Accessed 4 Feb 2019
3. Perry, J.R.: The future is in supercomputers. Whitehouse Art.: Infrastruct. Technol. 3 May 2018. <https://www.whitehouse.gov/articles/the-future-is-in-supercomputers/>. Accessed 4 Feb 2019
4. Simonite, T.: The US again has the world's most powerful supercomputer. Wired 8 June 2018. <https://www.wired.com/story/the-us-again-has-worlds-most-powerful-supercomputer/>. Accessed 4 Feb 2019
5. Lohr, S.: Move over, China: U.S. is again home to world's speediest supercomputer. The New York Times, 8 June 2018. <https://www.nytimes.com/2018/06/08/technology/supercomputer-china-us.html>. Accessed 4 Feb 2019
6. Walker, J.: The DoD office you've never heard of—and why that's about to change. Defense News, 1 January 2018. <https://www.defensenews.com/opinion/commentary/2018/01/11/the-dod-office-youve-never-heard-of-and-why-thats-about-to-change/>. Accessed 31 Jan 2019
7. Vanian, J.: Hewlett Packard enterprise signs big supercomputer deal with Department of Defense. Fortune.com, 15 February 2018. <http://fortune.com/2018/02/15/hewlett-packard-enterprise-supercomputer-department-defense/>. Accessed 31 Jan 2019
8. McCrystal, S., Collins, T., Silverman, D., Fussell, C.: Team of Teams: New Rules of Engagement for a Complex World. Penguin Random House LLC, New York (2015)
9. McInnes, L.: What Are Strategies for More Effective Teams? December 2017. <https://bssw.io/resources/what-are-strategies-for-more-effective-teams>. Accessed 4 Feb 2019
10. Moulton, D.J., Raybourn, E.M., McInnes, L., Heroux, M.: Enhancing Productivity and Innovation in ECP with a Team of Teams Approach. SAND2018-3987 C. figshare. <https://doi.org/10.6084/m9.figshare.6151097.v1> (2018)
11. Fussell, C., Goodyear, C.W.: One Mission: How Leaders Build a Team of Teams. Penguin Random House LLC, New York (2017)

12. Raybourn, E.M.: A new paradigm for serious games: transmedia learning for more effective training and education. *J. Comput. Sci.* **5**(3), 471–481 (2014). Elsevier
13. Raybourn, E.M.: Addressing changing mindsets: transforming next generation leader development with transmedia learning. In: Hailes, T.C., Wells, L. (eds.) *Changing Mindsets to Transform Security: Leader Development for an Unpredictable and Complex World*. National Defense University, Washington (2013)
14. Hansen, M.T.: *Great at Work: How Top Performers Do Less, Work Better, and Achieve More*. Simon & Schuster, New York (2018)
15. Mui, C.: Lead like a gardener. *Forbes* (2016). <https://www.forbes.com/sites/chunkamui/2016/07/14/to-nurture-innovation-lead-like-a-gardener-not-napoleon-or-kasparov/#5ef2ce7a3d83>. Accessed 4 Feb 2019