

Independent market research and
competitive analysis of next-generation
business and technology solutions for
service providers and vendors

**HEAVY
READING**
**WHITE
PAPER**

Heavy Reading's Open RAN Platforms and Architectures Operator Survey

*A Heavy Reading white paper produced for
AMD, Intel, and Kontron*



AUTHOR: SIMON STANLEY, ANALYST AT LARGE, HEAVY READING

INTRODUCTION AND KEY FINDINGS

Open RAN promises significant opportunities for carriers and enterprises to deliver 5G services without being tied to legacy equipment and a single supplier. Deployments are starting and are already capable of matching integrated RAN in simpler, lower end deployments. The current challenge is delivering the benefits of open RAN in high end deployment scenarios with large operating bandwidths and advanced antenna processing.

Open RAN refers to the ability to integrate, deploy, and operate radio access networks using components, subsystems, and software sourced from multiple suppliers. Standardized interfaces and functionality are now well defined by groups, including the O-RAN Alliance and the TIP OpenRAN project group, as well as 3GPP. Open RAN architectures build on developments for virtualized RAN (vRAN) and support the shift to open xHaul networks and edge computing. For vRAN, the networking functions are implemented in software and separated from the hardware. This approach allows for greater flexibility with the disaggregation of network functions and the use of commercial off-the-shelf (COTS) hardware.

Dramatic improvements in open RAN and vRAN network performance and bandwidths can be achieved by choosing the right RAN architecture and deploying server platforms that integrate the most appropriate combination of CPU processing and hardware acceleration. By choosing the best open RAN cloud and platform solutions, carriers and enterprises can deliver the wireless network and service performance expected by 5G users today, as well as provide the headroom needed to support developing standards and higher bandwidths as new spectrum and technologies become available.

This report presents the results of the Heavy Reading **Open RAN Platforms and Architectures Operator Survey** conducted during December 2022. The survey focused on investigating the opportunities and challenges in deploying open RAN/vRAN in operator networks and helping industry leaders and their customers evaluate the different approaches.

Key findings

- **The results suggest that most service providers believe the open RAN and vRAN ecosystem is ready for mass deployment now or will be by the end of 2024.** Of all respondents, 20% believe the open RAN/vRAN ecosystem is ready for mass deployment now, and a total of 71% believe it will be ready for mass deployment in the next 24 months.
- **The survey highlights that most respondents expect to have completed their first open RAN/vRAN installations in two years. However, there are significant differences between North America and Europe within those two years.** For Europe, 23% of respondents say they have already completed their first open RAN/vRAN installations, while for North America, just 8% of respondents say they have. For both North America and Europe, approximately 75% expect to have completed their first open RAN/vRAN installations in two years.


-
- **The results suggest private/enterprise networks and public networks are the highest priority for most service providers.** Of those responding, 71% say they are planning to prioritize the deployment of vRAN/open RAN for private/enterprise networks and 65% for public networks.
 - **The results indicate that open RAN will be widely deployed in both brownfield and greenfield, with a preference for greenfield deployments.** Of those responding, 40% will deploy open RAN in both brownfield and greenfield, whereas 33% will deploy open RAN in just greenfield, and 15% will deploy open RAN in just brownfield.
 - **Improved economies of scale for generic hardware and independent upgradability of hardware and software components are seen as the leading benefits of vRAN.** Of those responding, 35% see improved economies of scale for generic hardware as critical, and a further 60% see this as important. Additional benefits highlighted are improved power efficiency, network automation, rapid feature evolution, and the ability to leverage recent ecosystem innovations.
 - **A majority are planning to host artificial intelligence (AI), content/media delivery/hosting, cell site router (CSR), and firewall applications on their end open RAN/far edge solution, in addition to the base RAN stack.** More than 60% expect to support AI. Content/media delivery/hosting, CSR, and firewall are all in the top three other applications for at least 50% of respondents.
 - **The results suggest that operators are using a mix of hardware acceleration options, depending on the specific requirements of the RAN deployment.** Of those responding, 30% said they are already using both Layer 1(L1) lookaside and L1 inline distributed unit (DU) acceleration, and 80% are planning to use both in the next 24 months. A total of 36% are already using CPUs with integrated acceleration, and 82% are planning to use CPUs with integrated acceleration in the next 24 months.
 - **The results indicate that wireless vendors are using a variety of L1 baseband stacks, including those based on open-source software.** Of all respondents, 42% said their wireless vendors are using a third-party baseband stack, and 26% said their wireless vendors are using their own baseband stack. And of those responding, 27% said their wireless vendors are using an open-source baseband stack.
 - **Almost all respondents (97%) believe that operators do have a role in nurturing the open RAN ecosystem.** These included 44% that selected program management, 36% that selected systems integration, and 18% that selected solutions drivers for the role that operators should play in nurturing the open RAN ecosystem.

SURVEY DEMOGRAPHICS

The Heavy Reading **Open RAN Platforms and Architectures Operator Survey** was conducted online in December 2022. The questionnaire was written by Heavy Reading with input from project sponsors AMD, Intel, and Kontron. It was promoted to the Light Reading service provider database and received 113 responses from individuals working for operators with mobile network businesses. (Spurious, incomplete, and non-operator responses were removed.) All responses are confidential and are only presented in aggregate form. Heavy Reading does not share individual names or company names from the survey.

Figure 1 shows the responses by operator type. Of all respondents, 62% work for network operators, including 37% working for pure-play mobile operators and 25% for converged operators. The remaining respondents are split among enterprise service providers, cable operators with a mobile business, cloud or web services providers, and virtual operators.

Figure 1: What type of service provider do you work for?




n=113

Source: Heavy Reading

Note: Numbers in figures throughout this report may not total 100 due to rounding.

Figure 2 shows the response by geography. Western Europe is the largest market represented in the survey, with 30% of the response, followed by the US with 27%. Of all respondents, 50% work for organizations headquartered in Europe (Eastern and Western) and 32% in North America (US and Canada). This mix allows Heavy Reading to compare the responses from Europe and North America with reasonable confidence. Whenever the analysis in this report contrasts the largest and the smaller operators, it is made clear in the text.

Figure 2: In what region is your organization headquartered?




n=113

Source: Heavy Reading

Figure 3 shows responses by job title. Network operations and IT and cloud are the largest groups (61%), followed by R&D and technology strategy and corporate management.

Figure 3: What is your primary job function?




n=113

Source: Heavy Reading

Figure 4 shows the approximate annual revenue of the organizations covered by the survey. Of all respondents, 39% are from operators with at least \$1bn in annual revenue, with 18% from operators with revenue of more than \$5bn. The remaining 61% are from operators with less than \$1bn in annual revenue. This mix allows Heavy Reading to compare the responses from the larger operators (more than \$1bn) to those from the smaller operators with reasonable confidence. Whenever the analysis in this report contrasts the largest and the smaller operators, it is made clear in the text.

Figure 4: What are your organization's approximate annual revenues?



n=113


Source: Heavy Reading

OPEN RAN DEPLOYMENTS

For the first question, Heavy Reading asked when respondents believe the open RAN/vRAN ecosystem will be ready for mass deployment. The results shown in **Figure 5** suggest that a significant number of service providers believe that the open RAN/vRAN ecosystem is ready now for mass deployment, but the majority believe it will be ready for mass deployment by the end of 2024.

Of all respondents, 19% believe the open RAN/vRAN ecosystem is ready for mass deployment now, and a total of 38% believe it will be ready in the next 12 months. More than 70% believe the open RAN/vRAN ecosystem will be ready for mass deployment within the next two years.

Figure 5: When do you believe the open RAN/vRAN ecosystem will be ready for mass deployment?




n=113

Source: Heavy Reading

This confidence in the readiness of the ecosystem is also reflected in answers to the next question, with 15% of respondents saying that they have already completed their first open RAN/vRAN installations (see **Figure 6**). There are, however, some significant regional differences when these results are split out separately for North America (NA) and Europe. The results suggest that North America is moving slower than Europe but will catch up for the first installations by the end of 2024.

For Europe, 23% of respondents say they have already completed their first open RAN/vRAN installations, and more than half say they will have done so in the next 12 months. For North America, just 8% of respondents say they have already completed their first open RAN/vRAN installations, and 36% say they will have done so within the next 12 months. For both North America and Europe, approximately 75% expect to have completed their first open RAN/vRAN installations within two years.

Figure 6: When are you planning to complete your first open RAN/vRAN installations?



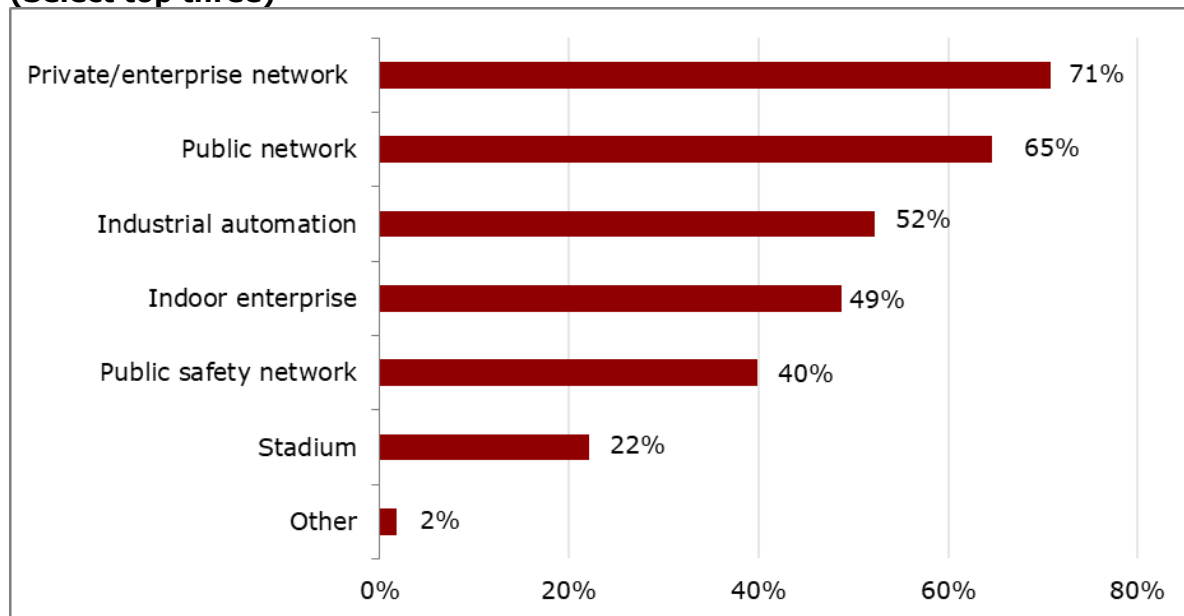
n=112

Source: Heavy Reading

The next few questions cover the types of deployment being planned and show the range of networks and environments that are supported by open RAN and vRAN implementations.

The results shown in **Figure 7** suggest that private/enterprise networks and public networks are the highest priority for most service providers. More than 70% say they are planning to prioritize the deployment of vRAN/open RAN for private/enterprise networks, and 65% say they are planning to prioritize the deployment of vRAN/open RAN for public networks. Industrial automation, indoor enterprise, and public safety networks are seen as less of a priority, with 40–52% saying they are planning to prioritize these deployments. Only 22% selected stadium deployments in their top three priorities.

Figure 7: Where do you plan to prioritize the deployment of vRAN/open RAN? (Select top three)



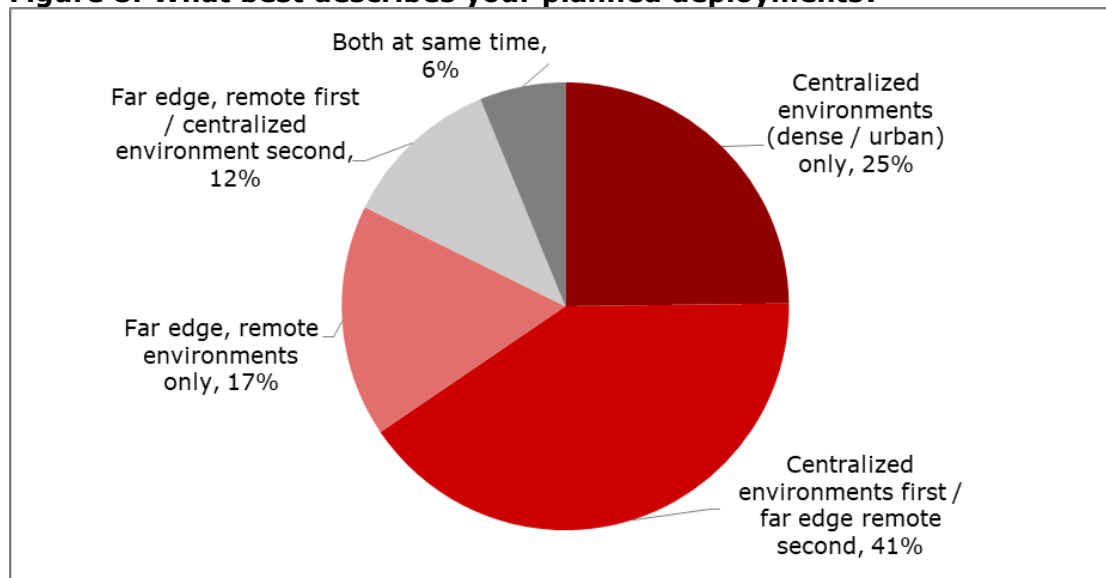
n=113

Source: Heavy Reading

Next, Heavy Reading asked about plans for deployments in dense/urban centralized environments and far edge, remote environments (see **Figure 8**). More than 70% are planning first deployments in dense/urban centralized environments, with the majority going on to deploy far edge, remote environments.

Of all respondents, 58% said they are planning a combination of centralized and far edge, remote deployments. One-quarter of respondents are planning just dense/urban centralized deployments, and 17% are planning just far edge, remote deployments. Of those responding, 12% will start with far edge, remote deployments before moving on to centralized deployments.

Figure 8: What best describes your planned deployments?



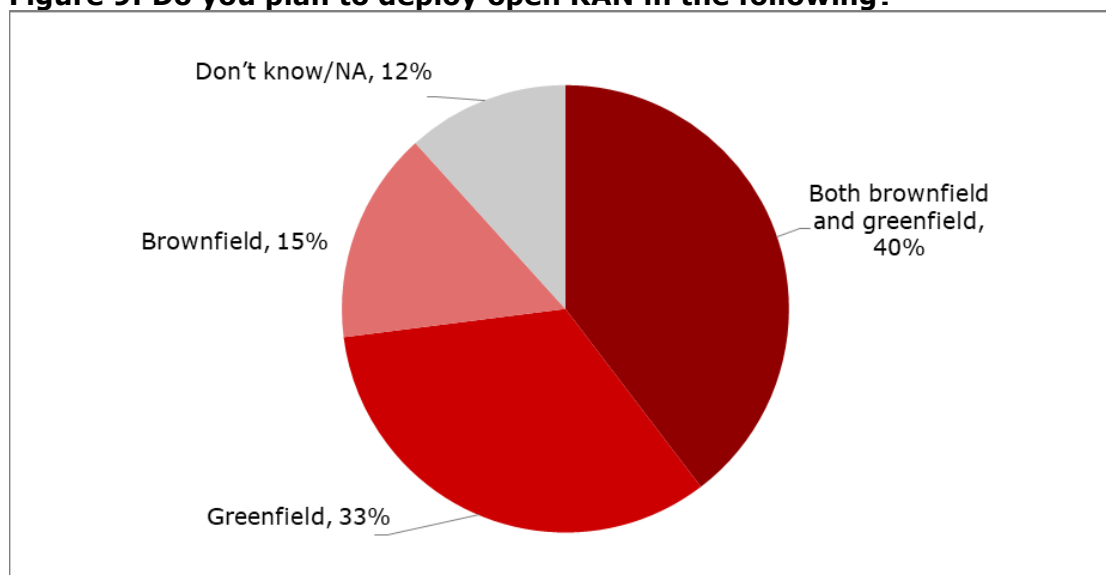
n=113

Source: Heavy Reading

Heavy Reading also asked about the split between greenfield and brownfield deployments. The results shown in **Figure 9** suggest that open RAN will be widely deployed in both brownfield and greenfield, with a preference for greenfield deployments.

Of all respondents, 40% will deploy open RAN in both brownfield and greenfield. One-third will deploy open RAN in just greenfield, and 15% will deploy open RAN in just brownfield.

Figure 9: Do you plan to deploy open RAN in the following?



n=111

Source: Heavy Reading


ARCHITECTURE AND EDGE CLOUD INTEGRATION

This next section looks at open RAN/vRAN architecture and edge cloud integration. For vRAN implementations, many of the virtualized network functions are running on server infrastructure located in edge data centers or cabinets. This edge cloud infrastructure also hosts other network functionality that benefits from being located close to the user, such as firewall, content delivery network caching and hosting, and AI to enhance network and application performance. Servers located in edge data centers or cabinets need to support a wide range of operating temperatures. These requirements are explored later in this section.

Figure 10 shows how respondents rank the main benefits of vRAN. The clear winner is the improved economies of scale for generic hardware, with 35% seeing this as critical and 95% seeing this as critical or important. The use of generic hardware based on COTS solutions also supports the independent upgradability of hardware and software components, which is seen by respondents as the second most important benefit of vRAN.

The use of edge cloud infrastructure to host both vRAN network functions and other network functions brings significant benefits from the pooling of resources, including generic hardware and cloud native software infrastructure. These benefits can include improved power efficiency, network automation, rapid feature evolution, and the ability to leverage recent ecosystem innovations. More than 80% of respondents see each of these benefits as critical or important.

Figure 10: How would you rank the main benefits of virtualized RAN?

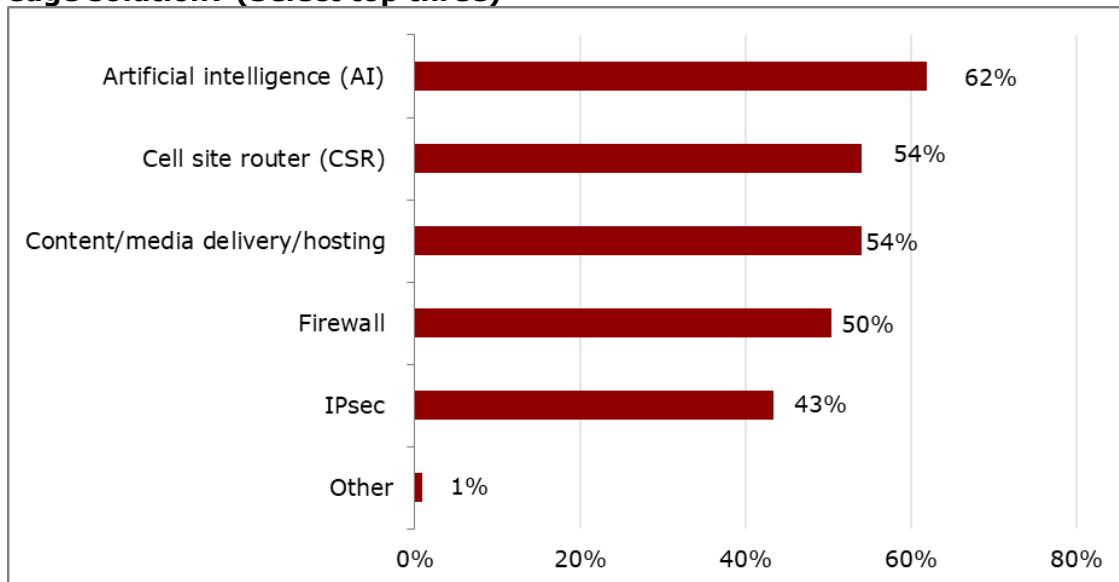


n=113

Source: Heavy Reading

Open RAN platforms and far edge servers can support multiple applications in addition to the base RAN protocol stack. For the next question, Heavy Reading asked which other applications respondents are planning to support (see **Figure 11**). More than 60% expect to support AI, and content/media delivery/hosting, CSR, and firewall are all in the top three other applications for at least 50% of respondents. IPsec is in the top three other applications at 43%.

Figure 11: In addition to the base RAN protocol stack to support wireless traffic, what other applications are you planning to support in your end open RAN/far edge solution? (Select top three)




n=113

Source: Heavy Reading

Low latency AI native applications are expected to have significant benefits when deployed within the centralized unit (CU) or DU solution. As shown in **Figure 12**, almost 60% expect to see benefits with network automation and optimization from deploying AI/machine learning (ML) in the CU/DU. Approximately 45% expect to see benefits with security/malware detection, power optimization, anomaly detection, or new revenue-generating business models.

Figure 12: Low latency AI native applications for CU/DU are beginning to be deployed. What benefits are you expecting from AI/ML in the CU/DU?




n=113

Source: Heavy Reading

The use of generic server hardware is key to delivering many of the benefits of open RAN and vRAN. These COTS platforms can be branded or white box server systems that are designed to meet service provider requirements for performance and operating environment. To find out how quickly operators are moving to these platforms, Heavy Reading asked when organizations expect parts of the RAN to move to a white box/COTS model.

The results in **Figure 13** suggest that the organizations covered by the survey are moving to a white box/COTS model of the radio unit (RU), DU, and CU at the same time, although the CU is slightly ahead for later deployments. Approximately 40% are already using a white box/COTS model for RU, DU, or CU or will do so within six months.

Figure 13: When does your organization expect parts of the RAN to move to a white box/COTS model?




n=113

Source: Heavy Reading

Many hardware platforms for open RAN and vRAN will need to support outdoor environments or indoor cabinets without controlled temperature. When asked about operating environments, almost 70% said they have outdoor applications (see **Figure 14**).

Of all respondents, 57% said they have indoor cabinets with controlled temperature, and the same number said they have indoor cabinets without controlled temperature. Fewer respondents said they had operating environments with minimum ambient temperature below -5°C (29%) or maximum ambient temperature above 55°C (18%).

Figure 14: Which best describes your operating environments?



n=113

Source: Heavy Reading

OPEN RAN/VRAN BASEBAND IMPLEMENTATION

This next section covers some of the options for implementing baseband processing in open RAN and vRAN deployments.

Open RAN and vRAN require L1 baseband processing within the DU that can handle the bandwidth and number of connections expected. Generic servers with high performance processors will support virtualized baseband implementations for limited bandwidth and number of connections. To support the higher bandwidths and the many connections required for more demanding wireless network deployments, such as those supporting macro cells, hardware acceleration is required.


There are three leading options for hardware acceleration to enhance DU performance on generic server hardware:

- **L1 Lookaside:** The L1 functions are running on the CPU, which offloads processing-intensive elements to a PCI card plugged into the server.
- **L1 Inline:** The full L1 functionality is implemented on a separate hardware system that may be a PCI card plugged into the server.
- **CPU with Integrated Acceleration:** The L1 functions are running on the CPU, which offloads processing-intensive elements to hardware engines integrated with the CPU.

Heavy Reading asked respondents what DU acceleration they are using for 5G deployments. The results shown in **Figure 15** suggest that operators are using a mix of hardware acceleration options depending on the specific requirements of the RAN deployment.

Of all respondents, 30% said they are already using both L1 lookaside and L1 inline DU acceleration, and 80% are planning to use both in the next 24 months. Of those responding, 36% are already using CPUs with integrated acceleration, and 82% are planning to use CPUs with integrated acceleration in the next 24 months.

Figure 15: What DU acceleration are you using for 5G deployments?



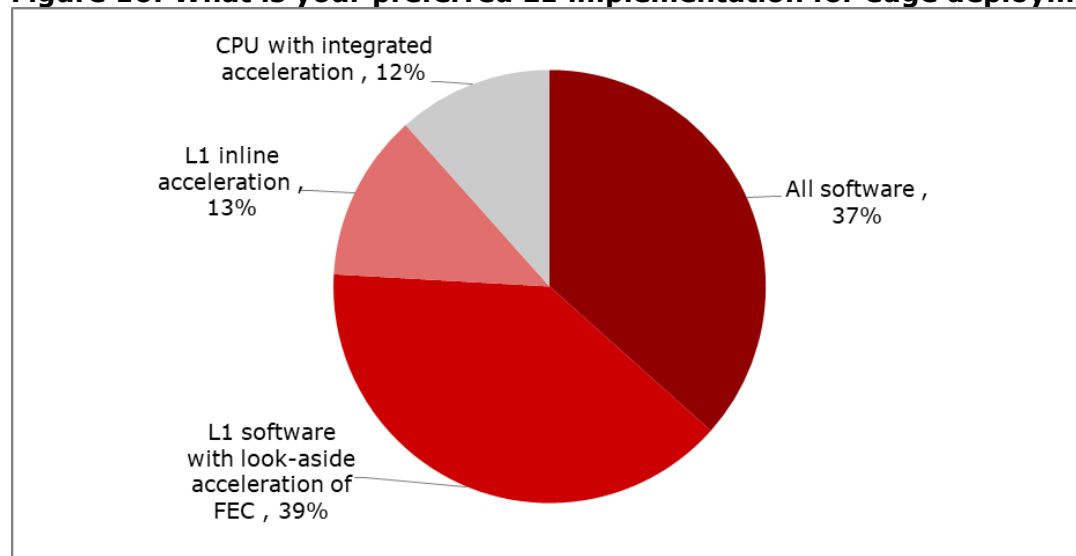
n=111

Source: Heavy Reading

Heavy Reading also asked about the preferred L1 implementations for edge deployments. The results shown in **Figure 16** suggest that most of the respondents are expecting server processors to provide adequate performance for edge deployments without hardware acceleration or with either lookaside or integrated acceleration.

For edge deployments, almost 40% said they prefer an L1 software implementation with lookaside acceleration, and a further 37% said they prefer an all-software implementation. Of all respondents, 13% prefer L1 inline acceleration for edge deployments, and 12% prefer a CPU with integrated acceleration implementation.

Figure 16: What is your preferred L1 implementation for edge deployments?




n=112

Source: Heavy Reading

One concern that has been expressed about using inline acceleration is that it could create a silicon vendor lock-in. The results shown in **Figure 17** suggest that this is a significant concern, with most respondents (74%) seeing inline acceleration as creating a silicon vendor lock-in. A significant minority (26%), however, do not see inline acceleration as creating a silicon vendor lock-in.

It is Heavy Reading's view that the risk of silicon vendor lock-in can be substantially reduced by ensuring that all hardware acceleration solutions are designed to work with open-source software and open RAN interfaces.

Figure 17: Do you see inline acceleration as creating a silicon vendor lock-in?



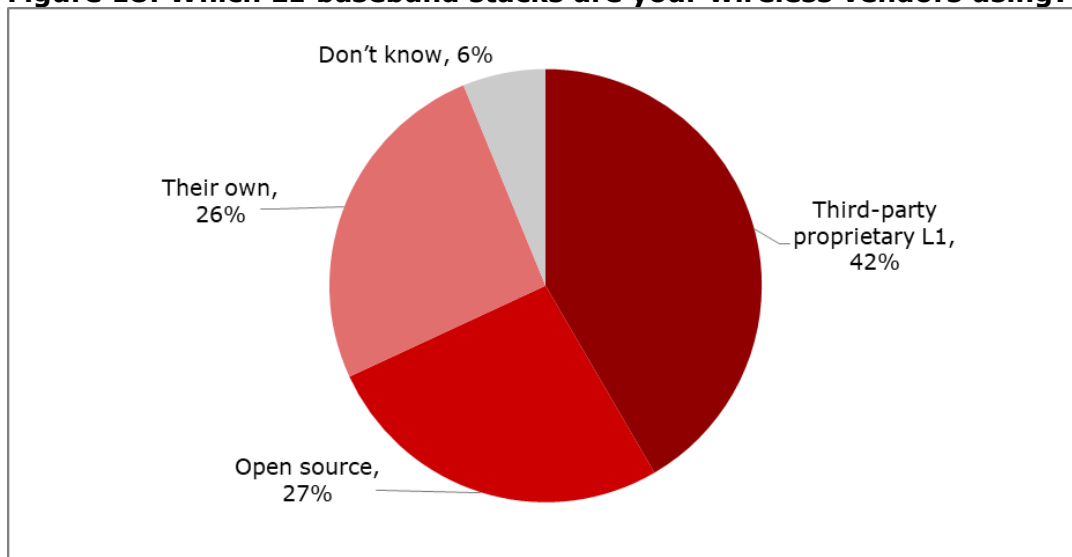
n=112

Source: Heavy Reading

For this next question, Heavy Reading asked respondents which L1 baseband stack their wireless vendors are using. The results shown in **Figure 18** suggest that wireless vendors are using a variety of L1 baseband stacks, including those based on open-source software.

Of all respondents, 42% said their wireless vendors are using a third-party baseband stack, and 26% said their wireless vendors are using their own baseband stack. Of those responding, 27% said their wireless vendors are using an open-source baseband stack.

Figure 18: Which L1 baseband stacks are your wireless vendors using?




n=113

Source: Heavy Reading

Cloud native software provides significant benefits over traditional virtualization approaches. Heavy Reading expects most vRAN and open RAN deployments to be using cloud native software in the next two to three years.

This view is confirmed by the results shown in **Figure 19**, which suggest that more than half of service providers will be deploying cloud native software for CU and DU in 24 months and that the vast majority will be deploying cloud native software for CU and DU in three years. Approximately 30% say they are already deploying cloud native software for CU or DU.

Figure 19: When will you be deploying cloud native software for CU and DU?



n=113


Source: Heavy Reading

KEY CONSIDERATIONS WHEN DEPLOYING OPEN RAN

This last section covers some key considerations when deploying open RAN and vRAN.

Heavy Reading wanted to understand the key areas that need improvement before carriers deploy open RAN/vRAN solutions at scale. The results shown in **Figure 20** show that improvements to the interoperability of hardware/software (HW/SW) components and the availability of qualified software vendors and systems integrators are important for most respondents. Both were selected by more than 60% of respondents. Also selected by more than 50% of respondents were security and hardware OEM vendor choice.

Figure 20: What areas need improvement prior to you deploying open RAN/vRAN solutions at scale? (Select all that apply)




n=113

Source: Heavy Reading

Heavy Reading also asked about the leading considerations for respondents when selecting an open RAN/vRAN platform solution. The results in **Figure 21** suggest that solution maturity, ease of use, and cost are the most important for carriers when selecting an open RAN/vRAN platform solution. Power consumption, supply, and feature completeness are seen as critical or important by almost 80% of respondents.

Figure 21: Rank the following items based on level of importance when selecting an open RAN/vRAN platform solution



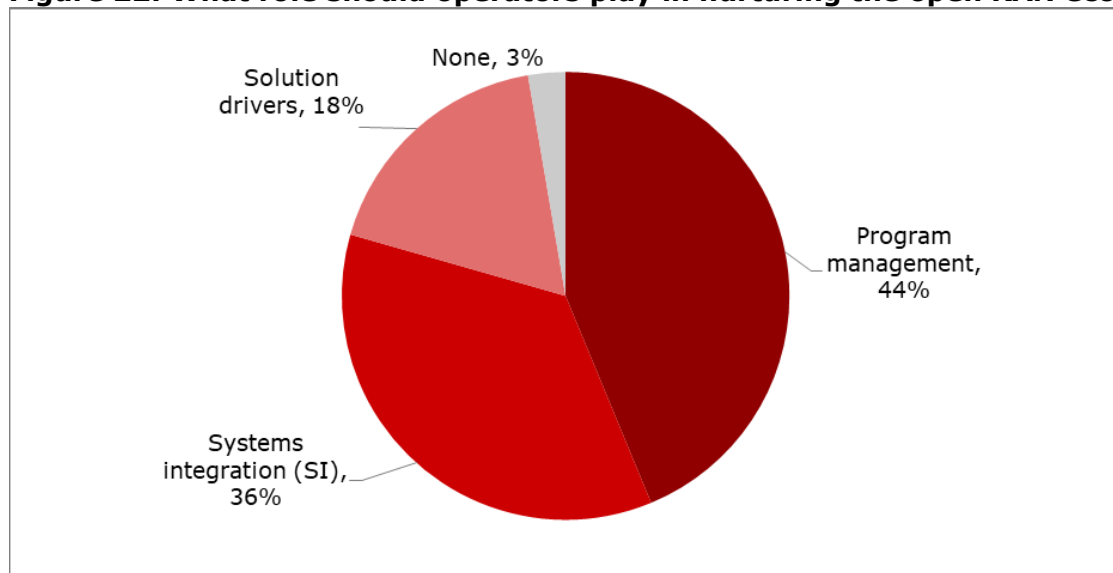
n=113

Source: Heavy Reading

The open RAN specifications have been largely completed over the last few years, but the ecosystem is still developing. Heavy Reading was keen to find out what role respondents believe operators have in nurturing the open RAN ecosystem.

The results in **Figure 22** show that 97% of respondents believe that operators do have a role in nurturing the open RAN ecosystem. These included 44% that selected program management and 36% that selected systems integration.

Figure 22: What role should operators play in nurturing the open RAN ecosystem?



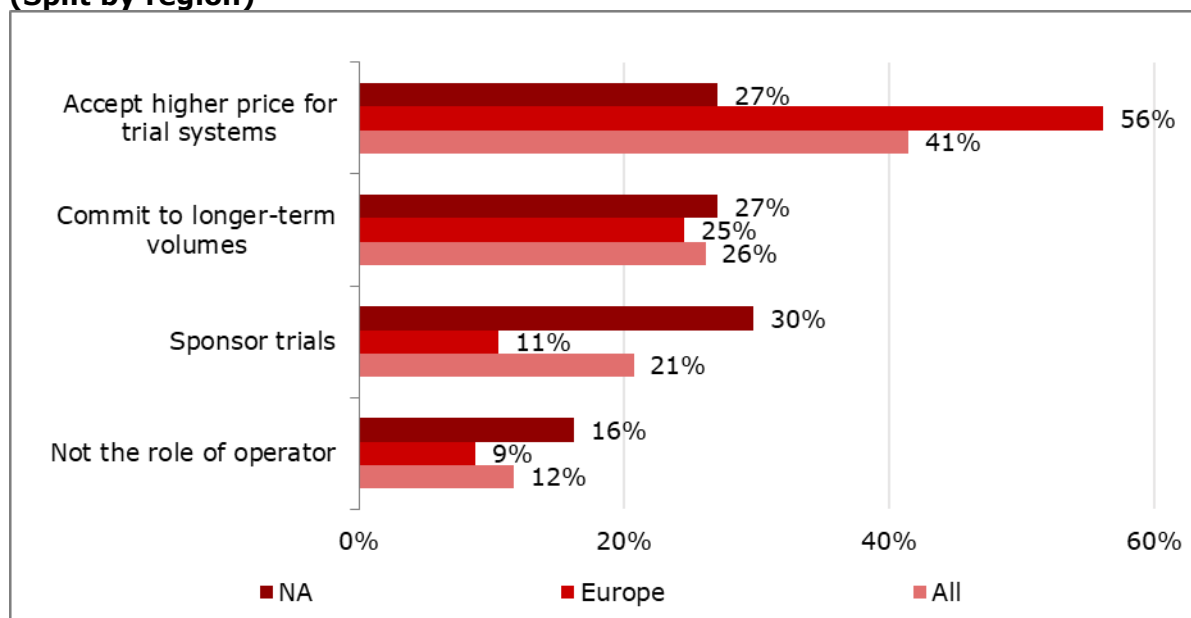
n=112

Source: Heavy Reading

For the open RAN ecosystem to continue developing, solutions developers need to see a path to revenue. Heavy Reading therefore asked respondents how operators can show a path to revenue for open RAN developers. The results are shown in the next two charts, split out by region and company size.

As shown in **Figure 23**, there were some significant regional differences when respondents were asked how operators can show a path to revenue for open RAN developers. From North America, approximately 30% of respondents selected sponsor trials, accept higher price for trial systems, and commit to longer-term volumes. From Europe, the majority (56%) selected accept higher price for trial systems, 25% selected commit to longer-term volumes, and just 11% selected sponsor trials.

Figure 23: How can operators show a path to revenue for open RAN developers? (Split by region)




n=111

Source: Heavy Reading

The results also show significant differences between large (more than \$1bn) and smaller operators (see **Figure 24**).

For larger operators, 49% selected accept high price for trial systems, and just 13% selected sponsor trials. For smaller operators, these options were evenly split, with approximately 30% selecting accept higher price for trial systems and sponsor trials. Among smaller operators, 18% said it was not the role of the operator to show a path to revenue for open RAN developers versus 7% for larger operators.

Figure 24: How can operators show a path to revenue for open RAN developers? (Split by company size)




n=111

Source: Heavy Reading

For the final question, Heavy Reading asked if respondents were concerned about the supply chain/source country for their open RAN vendor selection. The results shown in **Figure 25** suggest there is significant concern, with 77% saying they are concerned about the supply chain/source country for their open RAN vendor selection versus 23% that said they are not concerned.

Figure 25: Are you concerned about the supply chain/source country for your open RAN vendor selection?



n=111

Source: Heavy Reading