

Ericsson MSC Server Blade Cluster

The MSC-S Blade Cluster, the future-proof server part of Ericsson's Mobile Softswitch solution, provides very high capacity, effortless scalability, and outstanding system availability. It also means lower OPEX per subscriber, and sets the stage for business-efficient network solutions.

✧ PETRI MAEKINIEMI
✧ JAN SCHEURICH

The Mobile Switching Center Server (MSC-S) is a key part of Ericsson's Mobile Softswitch solution, controlling all circuit-switched call services, the user plane and media gateways. And now, with the MSC-S Blade Cluster, Ericsson has taken its Mobile Softswitch solution one step further, substantially increasing node availability and server capacity. The MSC-S Blade Cluster also dramatically simplifies the network, creating an infrastructure which is always available and easy to manage and which can be adjusted to handle increases in traffic and changing needs.

Continued strong growth in circuit-switched voice traffic necessitates fast and smooth increases in network capacity. The typical industry approach to increasing network capacity is to intro-

duce high-capacity servers – preferably scalable ones. And the historical solution to increasing the capacity of MSC servers has been to introduce a more powerful central processor. Ericsson's MSC-S Blade Cluster concept, by contrast, is an evolution of the MSC server architecture, where server functionality is implemented on several generic processor blades that work together in a group or cluster.

The very-high-capacity nodes that these clusters form require exceptional resilience both at the node and network level, a consideration that Ericsson has addressed in its new blade cluster concept.

Key benefits

The unique scalability of the MSC-S Blade Cluster gives network operators great flexibility when building their mobile softswitch networks and expanding network capacity as traffic grows, all without complicating network topology by adding more and more MSC servers.

The MSC-S Blade Cluster fulfills

the extreme in-service performance requirements put on large nodes. First, it is designed to ensure zero downtime – planned or unplanned. Second, it can be integrated into established and proven network-resilience concepts such as MSC in Pool.

Because MSC-S Blade Cluster operation and maintenance (O&M) does not depend on the number of blades, the scalability feature significantly reduces operating expenses per subscriber as node capacity expands.

Compared with traditional, non-scalable MSC servers, the MSC-S Blade Cluster has the potential to reduce power consumption by up to 60 percent and the physical footprint by up to 90 percent thanks especially to its optimized redundancy concepts and advanced components. Many of its generic components are used in other applications, such as IMS.

The MSC-S Blade Cluster also supports advanced, business-efficient network solutions, such as MSC-S nodes that cover multiple countries or are part of a shared mobile-softswitch core network (**Table 1**).

Key components

The key components of the MSC-S Blade Cluster (**Figure 1**) are the MSC-S blades, a signaling proxy (SPX), an IP load balancer, an I/O system and a site infrastructure support system (SIS).

The MSC-S blades are advanced generic processor boards, grouped into clusters and jointly running the MSC server application that controls circuit-switched calls and the mobile media gateways.

The signaling proxy (SPX) serves as

BOX A Terms and abbreviations

APG43	Adjunct Processor Group version 43	MSC	Mobile switching center
ATM	Asynchronous transfer mode	MSC-S	MSC server
E-GEM	Enhanced GEM	O&M	Operation and maintenance
GEM	Generic equipment magazine	OPEX	Operating expense
GEP	Generic processor board	OSS	Operations support system
IMS	IP Multimedia Subsystem	SIS	Site infrastructure support system
IO	Input/output	SPX	Signaling proxy
IP	Internet protocol	SS7	Signaling system no. 7
M-MGw	Media gateway for mobile networks	TDM	Time-division multiplexing

the network interface for SS7 signaling traffic over TDM, ATM and IP. It distributes external SS7 signaling traffic to the MSC-S blades. Two SPXs give 1+1 redundancy. Each SPX resides on a double-sided APZ processor.

The IP load balancer serves as the network interface for non-SS7-based IP-signaling traffic, such as the session initiation protocol (SIP). It distributes external IP signaling to the MSC-S blades. Two IP load-balancer boards give 1+1 redundancy.

The I/O system handles the transfer of data – for example, charging data, hot billing data, input via the man-machine interface, and statistics to and from MSC-S blades and SPXs. The MSC-S Blade Cluster I/O system uses the APG43 (Adjunct Processor Group version 43) and is 1+1 redundant. It is connected to the operation support system (OSS). Each I/O system resides on a double-sided processor.

The site infrastructure support system, which is also connected to the operation support system, provides the I/O system to Ericsson's Infrastructure components.

Key characteristics

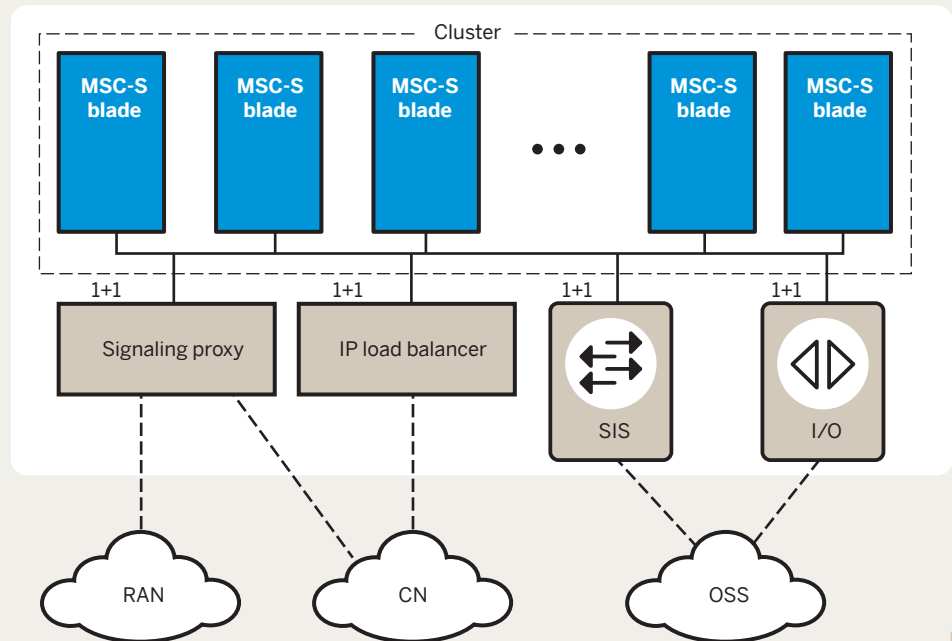
Compared with traditional MSC-S nodes, the MSC-S Blade Cluster offers breakthrough advances in terms of redundancy and scalability.

Redundancy

The MSC-S Blade Cluster features tailored redundancy schemes in different domains. The network signaling interfaces (C7 and non-C7) and the I/O systems work in a 1+1 redundancy configuration. If one of the components is unavailable, the other takes over, ensuring that service availability is not affected, regardless of whether component downtime is planned (for instance, for an upgrade or maintenance actions) or unplanned (hardware failure).

The redundancy scheme in the domain of the MSC-S blades is $n+1$. Although the individual blades do not feature hardware redundancy, the cluster of blades is fully redundant. All blades are equal, meaning that every blade can assume every role in the system. Furthermore, no subscriber record, mobile media gateway, or neighboring node is bound to any given blade.

FIGURE 1 Main functional components of the MSC-S Blade Cluster.



Therefore, the MSC-S Blade Cluster continues to offer full service availability when any of its blades is unavailable. Subscriber records are always maintained on two blades to ensure that subscriber data cannot be lost.

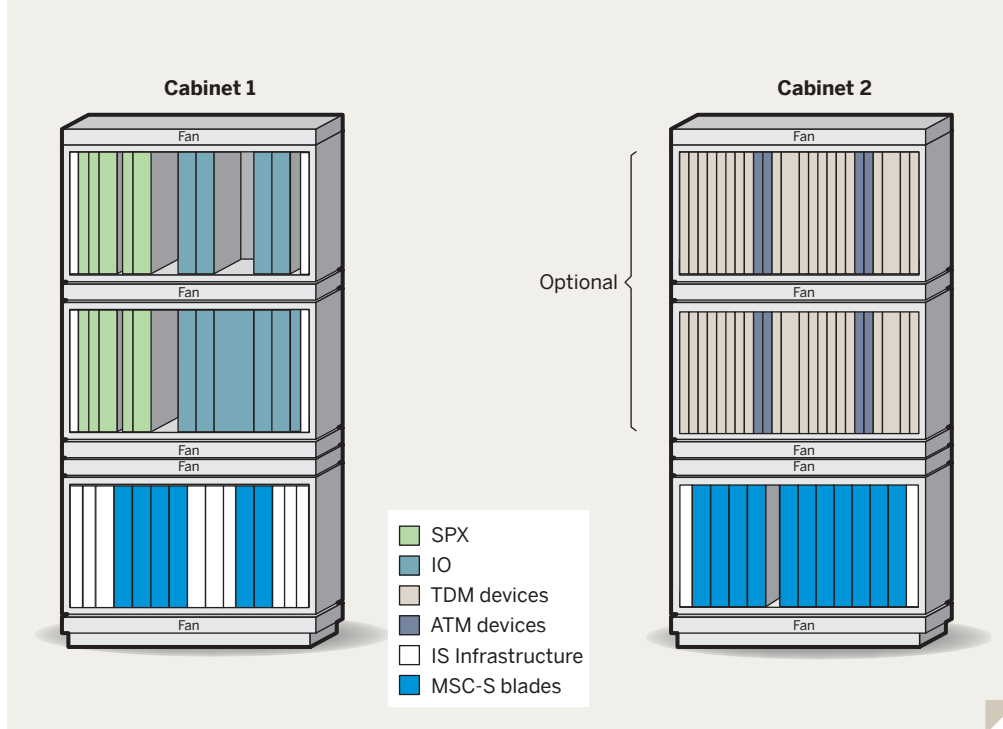
In the unlikely event of simultaneous failure of multiple blades, the cluster

will remain fully operational, experiencing only a loss of capacity (roughly proportional to the number of failed blades) and minor loss of subscriber records.

With the exception of very short interruptions that have no effect on in-service performance, blade

TABLE 1 Key benefits of the MSC-S Blade Cluster.

Feature	Benefit
Very high capacity	✦ Ten-fold increase over traditional MSC servers
Easy scalability	✦ Expansion in steps of 500,000 subscribers through the addition of new blades ✦ No network impact when adding blades
Outstanding system availability	✦ Zero downtime at the node level ✦ Upgrade without impact on traffic ✦ Network-level resilience through MSC in Pool
Reduced OPEX	✦ Fewer nodes or sites needed in the network ✦ Up to 90 percent smaller footprint ✦ As much as 60 percent lower power consumption
Business-efficient network solutions	✦ Multi-country operators ✦ Shared networks
Future-proof hardware	✦ Blades can be used in other applications, such as IMS

FIGURE 2 MSC-S Blade Cluster cabinets.

failures have no affect on connectivity to other nodes. Nor do blade failures affect availability for traffic of the user-plane resources controlled by the MSC-S Blade Cluster.

To achieve $n+1$ redundancy, the cluster of MSC-S blades employs a set of advanced distribution algorithms. Fault-tolerant middleware ensures that the blades share a consistent view of the cluster configuration at all times. The MSC-S application uses stateless distribution algorithms that rely on this cluster view. The middleware also provides a safe group-communication mechanism for the blades.

Static configuration data is replicated on every blade. This way, each blade that is to execute a requested service has access to the requisite data. Dynamic data, such as subscriber records or the state of external traffic devices, is replicated on two or more blades.

Interworking between the two redundancy domains in the MSC-S Blade Cluster is handled in an innovative manner. Components of the 1+1 redundancy domain do not require detailed information about the distribution of tasks and roles among the MSC-S blades.

The SPX and IP load balancer, for example, can base their forwarding decisions on stateless algorithms. The blades, on the other hand, use enhanced, industry-standard redundancy mechanisms when they interact with the 1+1 redundancy domain for, say, selecting the outgoing path.

The combination of cluster middleware, data replication and stateless distribution algorithms provides a distributed system architecture that is highly redundant and robust.

One particular benefit of $n+1$ redundancy is the potential to isolate an MSC-S blade from traffic to allow maintenance activities – for example, to update or upgrade software without disturbing cluster operation. This means zero planned cluster downtime.

Scalability

The MSC-S Blade Cluster was designed with scalability in mind: to increase system capacity one needs only add MSC-S blades to the cluster. The shared cluster components, such as the I/O system, the SPX and IP load balancer, have been designed and dimensioned to support

a wide range of cluster capacities from very small to very large.

The individual MSC-S blades are not visible to neighboring network nodes, such as the BSC, RNC, M-MGw, HLR, SCP, P-CSCF and so on. This first enabler is essential for smooth scalability: blades can be added immediately without affecting the configuration of cooperating nodes.

Other parts of the network might also have to be expanded to make full use of increased blade cluster capacity. When this is the case, these steps can be decoupled and taken independently.

The second enabler is the ability of the MSC-S Blade Cluster to dynamically adapt its internal distribution to a new blade configuration without manual intervention. As a consequence, the capacity-expansion procedure is almost fully automatic – only a few manual steps are needed to add a blade to the running system.

When a generic processor board is inserted and registered with the cluster middleware, the new blade is loaded with a 1-to-1 copy of the application software and a configuration of the active blades. The blade then joins the cluster and is prepared for manual test traffic. For the time being it remains isolated from regular traffic.

The blades automatically update their internal distribution tables to the new cluster configuration and replicate all necessary dynamic data, such as subscriber records, on the added blade. These activities run in the background and have no affect on cluster capacity or availability.

After a few minutes, when the internal preparations are complete and test results are satisfactory, the blade can be activated for traffic. From this point on, it handles its share of the cluster load and becomes an integral part of the cluster redundancy scheme.

MSC-S Blade Cluster hardware

Building practice

The MSC-S Blade Cluster is housed in an Enhanced Generic Equipment Magazine (E-GEM). Compared with the GEM, the E-GEM provides even more power per subrack and better cooling capabilities, which translates into a smaller footprint.

Generic processor board

The Generic Processor Board (GEP) used for the MSC-S blades is equipped with an x86 64-bit architecture processor. There are several variants of the equipped GEP board, all manufactured from the same printed circuit board. In addition, the GEP is used in a variety of configurations for several other components in the MSC-S Blade Cluster, namely the APG43, the SPX and SIS, and other application systems.

Infrastructure components

The infrastructure components provide layer-2 and layer-3 infrastructure for the blades, incorporating routers, SIS and switching components. They also provide the main on-site layer-2 protocol infrastructure for the MSC-S Blade Cluster. Ethernet is used on the backplane for signaling traffic.

Hardware layout

The MSC-S Blade Cluster consists of one or two cabinets (**Figure 2**). One cabinet houses mandatory subracks for the MSC-S blades with Infrastructure components, the APG43 and SPX. The other cabinet can be used to house a subrack for expanding the MSC-S blades and two subracks for TDM or ATM signaling interfaces.

Power consumption

Low power consumption is achieved by using advanced low-power processors in E-GEMs and GEP boards. The high subscriber capacity of the blade-cluster node means very low power consumption per subscriber.

Conclusion

The MSC-S Blade Cluster makes Ericsson's Mobile Softswitch solution even better – one that is easy to scale both in terms of capacity and functionality. It offers downtime-free MSC-S upgrades and updates, and outstanding node and network availability. What is more, the hardware can be reused in future node and network migrations.

The architecture, which is based on a cluster of blades, is aligned with Ericsson's Integrated Site concept and other components, such as the APG43 and SPX. O&M is supported by the OSS.

The MSC-S Blade Cluster distributes subscriber traffic between available

blades. One can add, isolate or remove blades without disturbing traffic. The system redistributes the subscribers and replicates subscriber data when the number of MSC-S blades changes. Cluster reconfiguration is an automatic procedure; moreover, the procedure is invisible to entities outside the node. ❖

Petri Maekiniemi



❖ is master strategic product manager for the MSC-S Blade Cluster. He joined Ericsson in Finland in 1985 but has worked at Ericsson Eurolab Aachen in Germany since 1991. Over the years he has served as system manager and, later, as product manager in several areas connected to the MSC, including subscriber services, charging and MSC pool. Petri holds a degree in telecommunications engineering from the Helsinki Institute of Technology, Finland.

Jan Scheurich



❖ is a principal systems designer. He joined Ericsson Eurolab Aachen in 1993 and has been employed since then as software designer, system manager, system tester and troubleshooter for various Ericsson products, ranging from B-ISDN over classical MSC, SGSN, MMS solutions, to MSS. He currently serves as system architect for the MSC-S Blade Cluster. Jan holds a degree in physics from Kiel University, Germany.

Acknowledgements

❖ The authors thank Joe Wilke whose comments and contributions have helped shape this article. Joe is a manager in R&D and has driven several technology-shift endeavors, including the MSC-S Blade Cluster.